Sustainable Intermediation: Using Market Design to Improve the Provision of Sanitation*

Jean-François Houde[1,5], Terence Johnson[3],
Molly Lipscomb[4], Laura Schechter[1]

December 3, 2024

Abstract

Many markets in developing countries, particularly those focused on public services, suffer from market failures caused by search frictions, externalities, and the exercise of market power. We use mechanism design tools to estimate the welfare gains from alternative market designs in the market for sanitation. We design and implement a centralized market with just-in-time procurement auctions over three years from 2013-2016. Using randomization in the auctions and the observed bidding behavior of the service operators, we measure the households' willingness-to-pay, and firms' marginal costs, outside options, information rents, and expected profit from entry. We estimate the extent to which centralized auctions can reduce prices in the market and increase the take-up of mechanized desludging services under various market designs.

*We thank Madelyn Bagby, Josh Deutschman, Jared Gars, Shoshana Griffith, Ahmadou Kandji, Nitin Krishnan, Sarah Nehrling, and Cheikh Samb for excellent research assistance and the Bill and Melinda Gates Foundation for generous funding for the project. Thanks to Mbaye Mbeguere and the members of the Senegal Office of Sanitation (ONAS) as well as Water and Sanitation for Africa for their collaboration on this project. We thank Radu Ban, Lori Beaman, Charlie Holt, Seema Jayachandran, Cynthia Kinnan, Robert Porter, Mar Reguant, seminar participants at the Econometrica Society North American Winter Meetings, Dartmouth, NBER-Market Design Fall Meetings, Northwestern University, Princeton University, University of Michigan, University of Virginia, UW-Madison, UW-Milwaukee, and the Washington-Area Development Conference for their comments.

Affiliations: [1]University of Wisconsin-Madison. [2]University of Virginia. [3] University of Virginia. [4]University of Wisconsin-Madison. [5]NBER.

Developing country markets are often marked by factors that make them susceptible to high prices stemming from inefficiency and lack of competition. Many markets require high fixed-cost investments in an environment where credit can be difficult to procure, the number of service providers in the market is relatively small, the products are homogeneous, and government and large corporate contracts are the most lucrative components of the business.

We study one such market that is plagued with important market imperfections: the market for residential *desludging* in Dakar, Senegal. When a household latrine pit fills, the household must empty it by having it desludged. Households then have a choice between a "manual" option performed by a hired worker or family member, and a "mechanized" option that uses a vacuum truck. In most neighborhoods of the city, high mechanized desludging prices have resulted in the widespread use of manual desludging: twenty-five percent of households in our sample hire a manual worker to remove human fecal sludge from the premises, and 27% of households do it themselves–both result in disposal of the sludge in the immediate vicinity of the household (typically in the road in front of the house). The large externalities associated with sanitation services make the low equilibrium levels of sanitation provision particularly important: there is a high correlation between diarrheal disease and poor desludging practices (Deutschman et al, 2020) as well as sanitation and health in general (Fewtrell et al, 2005). Prior work suggests that the weflare effect of lowering prices can be quite large: the estimated elasticity of demand for mechanized desludging in Dakar is -1.8 (Deutschman et al, 2020).

The traditional market for mechanized desludging services in Senegal is a decentralized search market composed of independent and integrated operators. Once they have located an operator, households know that they would have to pay an additional search cost in order to re-enter the market if they are unable to come to an agreement on prices with the operator they have. Even once the terms of a job seem settled, an operator can always renege once they arrive at the household's premises. Moreover, allegations of tacit collusion between providers further limit the gain from search and reduce the supply of mechanized desludging.

In this paper, we quantify the benefits for consumers of an auction-based intermediation platform designed to increase competition and improve the match between households and service providers. In collaboration with the Senegalese National Office of Sanitation (ONAS), we instituted just-in-time procurement auctions for desludging services in Dakar, Senegal. Over the course of three years, from 2013 to 2016, over 5,000 clients searching for mechanized desludging providers contacted the call-center to obtain a price quote. On the supply-side, the platform randomized the number of invited bidders assigned to each client, as well as the auction format (revisable vs sealed-bid). The goal of this experimental platform was to identify the competitive conduct of providers and measure clients' willingness to pay. In a companion paper, Houde, Johnson, Limscomb, and Schechter (2024), we exploit this source experimental variation to test for tacit collusion in the market, and estimate the share of providers behaving competitively on the platform.

Our objective is to counter-factually predict the equilibrium effect of designing a competitive intermediation platform on consumer surplus and on the takeup of mechanized desludging. We do so first by estimating the fundamental parameters of operator costs and clients' demand using the random variation created by the experimental platform. Using a model of providers' participation and bidding, we then compare the observed platform take-up rates and costs with platform outcomes when the market

is managed under first-best marginal cost pricing and second-best auctions with optimized invitations.

In the process of designing and running the auctions, we found that there are many elements that have strong impacts on prices and take-up. First, we document that most invited bidders rely on non-competitive bidding strategies, characterized by a handful of focal prices and coarse bidding. Second, desludger bids and participation decisions vary substantially with distance and idiosyncratic differences, suggesting that invitation rules and asymmetric information play important roles. Third, bidders may enter strategically in order to maintain larger margins. Because they have more information than the planner over their outside options, bidders may not enter as frequently as a social planner maximizing consumer welfare would require.

These market features explain the relatively poor performance of the experimental platform. By randomly inviting providers (many with very low participation rates), the platform failed to invite a sufficiently large number of competitive bidders to each auction (i.e. bidders with low expected costs, and bidders willing to deviate from collusive prices). These design features implied that the distribution of quotes obtained from the call-center is comparable to the one obtained by consumers in the decentralized market, and less than 30% of quotes are accepted leading to a modest decrease in transaction prices. Moreover, nearly 20% of auctions failed to receive any bids. Despite this, we find that a significant share of suppliers learned quickly to bid strategically on the platform, and roughly 25% of active bidders compete aggressively for clients. This suggests eliminating the experimental elements in the platform and better designing the auctions could further lower prices for consumers.

This paper uses tools from mechanism design and industrial organization to quantify the benefit of improving the design of the auction platform. We focus, in particular, on measuring the equilibrium effect of targeting invitations on the auction platform and eliminating non-competitive conduct. To do so, we proceed in three steps.

We first exploit randomization in the number and characteristics of invited bidders to estimate the demand of consumers on the platform. We find that those random features induce substantial random variation in prices offered to consumers, and cannot reject the null hypothesis that winning bids are independent of unobserved heterogeneity in willingness-to-pay.[1] As a result, we estimate the distribution of willingness-to-pay using a simple Probit model. Clients on the platform have a very steep demand curve: the average elasticity of demand evaluated at the observed price quotes is -3.5.

Next, we use desludger bids in the sample of first-price sealed-bid auctions to estimate costs by modifying the Guerre, Perrigne, and Vuong (2000) framework to account for ties and discrete bids. In particular, we account for differences in competitive conduct across bidders by inferring the consideration set of active bidders. Competitive types rely on a relatively fine price grid, which allows them to under-cut commonly use focal prices. Collusive and/or inattentive bidders, on the other hand, rely on a handful of prices, which soften price competition on the platform and lead to frequent ties. We assume that bidders maximize the expected profits of participating in each auction by selecting a bid within their respective consideration sets. This allows us to estimate a semi-parametric distribution of bidders' costs using an ordered-choice model. We find that the average marginal cost depends on the proximity between clients and firms' garages, leading to asymmetries between firms within the auctions. We also

[1] We construct instruments and measure demand using an instrumental variables probit model pioneered by Rivers and Vuong (1988).

find substantial residual dispersion in cost, consistent with the importance of asymmetric information on the platform.

Finally, we estimate the distribution of outside options by rationalizing observed participation with a sequential model of firm entry. As in Krasnokutskaya and Seim (2011), we assume that the opportunity cost of participating in the auction is conditionally independent of the private-value cost shock of providing the service. This allows us to model the strategic decision to bid as an entry game with asymmetric information. The results confirm that the expected profit of competing in the auction is the main determinant of participation and that dispersion in privation information about the opportunity cost of bidding is large (on par with the standard deviation of expected profits). This unobserved heterogeneity captures the fact that invitations are not targeting bidders who are necessarily available to provide the service, leading to uncertainty in the number of bids received.

With the fundamentals of the market recovered, we adopt a mechanism design approach. The Revelation Principle guarantees that for any equilibrium in a game of incomplete information, there is an alternative game called a Direct Revelation Mechanism in which agents report their private information as strategies and receive the same payoffs. As in Myerson (1981), we exploit this result to reduce the market design problem to a pair of simple questions: Who should be invited to a particular auction? Given their cost shocks, who should be selected to complete the job? By abstracting away from the particular rules of the auction — second-price, paid-as-bid, reserve prices, and so on — we can answer a number of important questions about the market: how much welfare is lost due to private information and strategic entry? How much could the performance of the platform be improved? What are the weaknesses of the current platform that reduce its performance?

The results can be summarized as follows. We first compare outcomes from the current platform by simulating a counter-factual equilibrium that holds fixed the invitation rules. In particular, we use the same uniform invitation probability and the same average number of bidders invited among those who submitted at least one bid. We find that eliminating non-competitive conduct on the platform substantially increases the number of bids received and lowers the average lowest bid received by roughly 7%. Since bids are more dispersed when we remove focal prices, the effect on transaction prices is larger ($\approx$ 15%), and the transaction rate would increase by nearly 10 percentage points (from 27% to 36%). In Houde et al. (2024) we constructed a similar counter-factual exercise by simulating the outcome of auctions in which all active bidders used the competitive-type policy function. This exercise ignores the equilibrium response of bidders from a market-wide change in conduct and as a result underestimates the gain from eliminating non-competitive conduct on demand (i.e. increases of 23% vs 35%).

Next, we analyze the effect of changing the platform invitation rules. Holding fixed the number of invited bidders, we find that targeting invitations by distance has a relatively small effect on the performance of the platform. An invitation rule that invites potential bidders using an inverse-distance probability weight increases the participation rate by one percentage point, and lowers the winning bid by 2%, and increases the take-up rate by 3.3% (from 35.9% to 37.1%).

On the other hand, increasing the number of bidders invited to each auction has a large effect on the average winning bid and the take-up rate. An invitation rule that invites 20 truckers based on distance (instead of 10) would increase the take-up rate from 37.1% to 42%. This increase in competition reduces the participation rate from 46.6% to 36.4% in our baseline specification (since expected profits

are lower), but a 7% reduction in the average winning bid outweighs this adverse effect. We find that further increasing the number of invited bidders to include all active bidders (39) would continue to have a pro-competitive effect. In our baseline specification, we find that this "optimal" platform would increase the take-up rate to 46%, and lower the average accepted offer to CFA 19,000. In contrast, we estimate from a large household survey that the average transaction price in the market is CFA 25,000, and the average accepted price in the experimental platform is CFA 24,100.

Finally, we can decompose the losses from private information and strategic entry compared to the hypothetical first-best outcome corresponding to perfect competition, where all desludgers bid their true costs. In particular, we estimate that the cost of private information amounts to a -23.1% loss in consumer surplus relative to the first-best, while strategic entry and private information together reduce consumer welfare -47%. The first-best allocation is associated with a nearly 70% take-up rate, more than double what we currently observe in the experimental platform.

We conclude that a platform that maximizes consumer surplus subject to incentive-compatible reporting of costs by firms and their participation constraints requiring that they expect entry to be profitable would generate 53% of the consumer surplus achievable in the first-best. Although these are substantial losses, the current platform is estimated to generate only 26.6% of the first-best consumer surplus. This illustrates the potential benefit of using a competitive auction platform to intermediate trade in this market.

The rest of the paper is organized as follows. In section 1 we discuss the sources of data for this paper, as well as our intervention and the market for sanitation services in Dakar, Senegal. Section 2 describes the empirical bidding model that we use, and our identification and estimation strategy. Section 3 presents the estimation results of the model primitives. Section 4 introduces the direct mechanism approach that we use to predict counterfactual outcomes and presents our main simulation results. In section 5 we discuss policy implications and conclude.

1 Market and data description

1.1 Demand and supply for desludging services

When a household latrine pit fills, the household must empty it by having it desludged. This tends to happen once or twice a year for households in Dakar. Households then have a choice between three types of desludging services: (i) manual performed by a family member, (ii) manual performed by a hired worker (called a "baay pell" in Senegal), and (iii) mechanized using a vacuum truck. Manual desludging consists of removing the sludge using a shovel and placing it in a pit dug in the street near the house.[2] Mechanized desludgings are done by two to three workers with a vacuum truck. The truck pumps as much sludge out of the pit as possible and takes it relatively far from the household, either dumping it legally at a treatment center or illegally in a street drainage canal or the ocean.

The market for mechanized desludging is organized around three treatment centers scattered across the city and a network of garages (or parking lots) where clients meet service providers. Based on

[2]The manual option increases the risk of health-related sanitation problems (for the client, the workers, and other households in the neighborhood). Manual desludging is technically illegal, though it is rarely sanctioned. It is often a source of controversy among neighbors.

our discussions with market participants, prices in the traditional market are determined by bilateral negotiation between clients and truck drivers. Consumers meet providers at garages, and haggle over prices. Walk-in clients at a parking lot are allocated to the driver who is first in line, and drivers from the same garage do not compete over clients (similar to what often happens at taxi stands). The proximity of consumers to garages therefore affects both their ability to search and to negotiate for better prices, as well the cost of providing the service. Furthermore, repeated interactions between truckers within (and potentially across) garages limit the bargaining leverage of consumers, and firms typically initiate the negotiation by relying on "focal" prices that are common across garages.

The rising popularity of cell phones in the years leading to our sample period have somewhat eroded the bargaining power of truckers. Since households need to desludge multiple times, many clients contact truck drivers directly; either by calling their phone number or hailing them on the street. This has allowed independent truckers to compete with companies affiliated with the main garages, and offer better prices.

This search and price negotiation process describes the functioning on the main peri-urban residential neighborhoods of Dakar in which most multi-truck companies are affiliated with the Association of Desludging Operators (or AAAS), which controls the operation and prices of the larger garages.[3] Firms working in the arrondissement of Rufisque on the eastern outskirts of Dakar which the Association does not control operates independently on the Association, and providers rely mostly on a posted-price strategy with three tiers (based on load size). Prices as a result are much lower in the arrondissement of Rufisque on the eastern outskirts of Dakar which the Association does not control.[4] Most of our empirical analysis focuses auctions performed outside of Rufisque.

The cost of providing a mechanized desludging includes the time and fuel required to complete the job (while pumping sludge at the client's house and driving from the garage to the client to the treatment center and back to the garage again), the treatment center's dumping fee,[5] and in some cases, a referral commission paid to the garage. Mechanized desludging exhibits economies of scale due to truck maintenance and/or rental costs. The large majority of drivers, 92%, are paid a fixed salary, and about half report paying a commission for jobs the garage refers to them. A portion of these revenues is redistributed to company owners in the form of revenue-sharing agreements. For example, desludgers in the largest garages report being paid by their garage on days when they do not find work.

Since driving the truck and operating the vacuum pump require significant amounts of fuel, the cost of diesel and the efficiency of the truck play an important role in determining trucker costs. The fuel efficiency of the trucks varies substantially with truck size and age, and we estimate that most trucks get between three and six kilometers per liter of diesel. Since diesel prices averaged 750 CFA per liter over

[3]We did not survey in the historical center of the city found towards the west of the map because it is relatively well connected to the sewer network. We did survey in the arrondissement of Rufisque, which is outlined in yellow on the east of the map. Rufisque is not controlled by the Association and has a more competitive mechanized desludging market.

[4]In the 1990s a single company, UPAMA, provided desludging services in the independent municipality of Rufisque. UPAMA's main line of business was desludging for fish product processing companies, and it was asked by the city administration to provide affordable residential desludging. Over time, new companies entered the market to serve growing demand, but the new companies matched UPAMA's base price in order to get business. The proximity of the Rufisque treatment plant helps reduce the variable cost of the service relative to other areas of Dakar. UPAMA receives no direct subsidies, and we believe the price reflects the cost of providing a mechanized desludging in the area.

[5]The fee for disposal of the sludge at a treatment center is approximately 3,000 CFA.

the period, the fuel cost per kilometer ranges between 125 CFA and 250 CFA.[6]

The decision of drivers to accept new contract is driven in part by capacity constraints. Conversations with market participants also reveal that a single latrine pit usually fills the truck more than halfway, limiting the ability of drivers to service multiple clients without dumping the sludge at a treatment center in between. About 8% of households require more than one trip. To estimate the production capacity of trucks, we ask operators how many desludging trips they performed over the last ten days. On average, trucks perform slightly more than one trip per day (residential and non-residential combined), but there is substantial heterogeneity across trucks. The top ten percent of trucks in terms of number of trips perform more than three jobs per day, and the most active truck performed over 50 trips in a ten-day period. Based on these data and discussions with providers, we estimate that a typical job takes about two hours from start to finish. Most truckers operate with substantial excess capacity, while only a small fraction operate at full capacity.

1.2 Intervention: Auction platform

Auctions are a theoretically and practically compelling method of minimizing procurement costs in the presence of uncertainty about supply. In order to increase competition in the market and reduce costs, we designed just-in-time auctions for procurement of desludging jobs to be run through a call center. Together with Water and Sanitation for Africa (WSA) and the National Office of Sanitation in Senegal (ONAS), the call center ran from July 2013-September 2015. The call center has since been scaled up by ONAS, was advertised across Dakar through billboards, radio ads and television ads, and was more recently taken over by a private company which runs the sewage treatment centers, Delvic. Call center activity is clustered around the peri-urban areas of Dakar, particularly Pikine and Guediawaye, but calls have come from downtown Dakar as well as further out of the city. Delvic, the current manager of the center, continues to share their data on more recent auctions with us and they are interested in understanding how to increase competition in the auctions.

When a household needs a desludging service, it calls the center and gives the dispatcher basic information about the location of the pit to be desludged. Jobs are made as homogeneous as possible through the rules of the call center: Households are told to expect the desludger to remove up to 8 cubic meters of sludge, and the winning operator is required to be available to do the job within two hours of the end of the auction.[7]

The call center dispatcher sends the job out for bidding to 8-20 randomly selected desludging operators via text message. Invitations are randomized between the 104 desludging operators in the system, but only approximately 43 participate consistently. Both independent and affiliated operators are invited to bid on every job, and bidding is done anonymously. This randomization of invitations allows for exogenous variation in the competition that a bidder faces across auctions. The platform informs the

[6] In order to estimate the amount of diesel necessary per kilometer for a job, we sent an enumerator on ride-alongs with two truck drivers, filling the tank at the beginning and end of the day and recording the kilometers traveled and diesel used.

[7] These rules were made so that desludgers could form expectations on the requirements of the service before submitting a bid and so that households could expect quick service when purchasing through the center. Desludgers with trucks less than 8 cubic meters are allowed to participate, but they must be ready to do more than one trip to empty a household's tank. Similarly, desludgers with tanks larger than 8 cubic meters are not required to take more than 8 cubic meters when they perform the service.

bidder as to the number of operators who have been invited and the number that are single truck operators versus those affiliated with larger companies, but not the specific identities of the other bidders. The bidder will know if other operators from their parking lots are invited (because they will all receive text messages at the same time), but they will not know which other operators who do not happen to be at the parking lot are also receiving the invitation.

The invited desludging operators have one hour to bid on the job under one of our two auction formats using text messages, and the operator with the lowest offer wins the job. In case of a tie, the bidder who sent the earlier bid wins, encouraging prompt bidding. The dispatcher then calls the household, informs them of the winning bid, and asks if the household accepts that price. If the household accepts, then the desludging operator is given the household-specific information, and they plan the remainder of the transaction directly with the household. The final winning bid and whether the winner was an independent desludger or a company is sent to all desludgers who have been invited to participate in the auction. The desludger who won the auction is alerted that they won, and subsequently receives a message with either the contact information of the household or a message letting them know that the household turned down the price. The winning desludger is then required to arrive at the household to do the job within two hours of winning the auction, and the household pays the desludger directly.

There are two auction formats, and which auction format is chosen for a particular auction is randomized with a 50% probability of each. In the first, the "closed format," the bidders receive no updates after the auction begins and have one hour to submit a single bid. Bidders in the closed format auctions are not allowed to change their bids after they have been submitted. In this analysis we focus on the closed format auctions. In the second, the "open format," the bidders are given updates about the standing low bid every 15 minutes and 5 minutes before the auction closes, and can submit a new, lower bid any time. In both cases, desludgers receive text messages reminding them that they have been invited to the auction at the same interval (every 15 minutes and 5 minutes before the end of the auction).

There are therefore three main layers of randomization in the auction platform from which we are able to derive unbiased estimates of the behavior of the desludgers: 1. the auction format, 2. the number of bidders invited to the auction, and 3. the identities of the bidders invited to the auction. We use this randomization in order to observe how the desludgers change their bidding strategies when they have more linked desludgers with them in the auctions, and when they are able to observe more of the action in the bidding process through the revisable bid auction format.

1.3 Data

The intervention and data collection spanned mid-2012 through mid-2015 in residential neighborhoods surrounding Dakar, Senegal. In this paper we use primarily two sources of data.

Our main data-set is a panel data from the just-in-time auction platform for the procurement of residential desludging jobs. The administrative data from these auctions includes the randomized auction format, the randomized number and identities of the desludgers invited, whether or not they bid, the time and amount of their bid if they made one, the location of the household that they were bidding on, and the winning bid. We use this information to construct characteristics of consumers that enter our demand and bidding model. We use the location of household to measure the distance to the nearest treatment (cost shifter) as well as indicator variables for one of the five main *Arrondissements* in Dakar.

We ran over 5,000 auctions for mechanized desludging services through the call center in collaboration with the Senegalese Office of Sanitation (ONAS) from July 2013 through October 28th, 2016.[8] Some auctions were performed for subsidized households (671 auctions), which affected the price paid by households, though not the price received by desludgers, and so these auctions are excluded from regressions studying household acceptance.[9] The sample sizes vary across auction-level regressions as follows. There are 5,162 auctions total; 4,177 of those have at least one bid; 2,865 of those have at least two bids; 4,491 are unsubsidized auctions; and 3,621 are unsubsidized auctions with at least one bid. The research team designed and managed the platform from its inception in 2013. In January 2015, the management of the platform was transferred to ONAS. Penalties were abolished and the invitation rule was modified so that all truckers had the same probability of being invited rather than over-sampling active participants. Also, the number of truckers invited decreased on average from 14 to 11. Although the changes to the invitation rule affected the performance of the platform (by reducing competition), the assignment of bidders to formats and auctions remained random. Slightly more than half of auctions were performed prior to the design change (55%). Because of the design change affected the distribution of bidders on the platform, most of our empirical analysis focuses on sealed-bid auctions conducted under the "old" plaform.

Table 1 presents summary statistics regarding the auctions and bidders. Around half of the auctions are randomized to be revisable format and half are sealed-bid. An average of 13 bidders are invited to each auction, and bidders were invited to bid in 524 auctions on average. The participation rate is fairly low. The probability of submitting a bid is about 10%, which leads to an average of 2.4 valid bids per auction, conditional on there being at least one bid. This low participation rate is due to the fact that a majority of desludgers rarely or ever bid, and the invitations were not targeted based on bidder availability or distance to client. As we discuss below, 52 bidders submit more than 20 bids during our sample periods and much of our bidder-level analysis focuses on the 39 most active bidders submitting more than 30 bids. Active bidders submit bids 24% of the time.

Our second source of data is a survey of desludging provider. We conducted a baseline survey of 121 desludging truck operators in mid-2012, and an endline survey of 152 drivers (of which 13 were owner-operators), 75 truck owners, and 20 managers in mid-2015. Our goal was to conduct a census of trucks active in the residential desludging market. An operator is either the manager of a fleet of trucks, or a driver associated with a single license plate.

We use this data to construct characteristics of potential bidders that affect the cost of participating and providing the service. Our main cost shifter is the distance between invited bidders and consumers. We measure distance using the location of the trucker's garage and the location of the client, and add the distance from the client to the nearest treatment center to estimate the total travel distance necessary to complete the job.

We use the results of the survey to measure the truck size, and an indicator variable equal to one if the bidder operates as a single-truck fleet. We also include an indicator variable equal to one if the

[8] Before the auctions began, the project held multiple training sessions for the truckers to help them understand the auction process and teach them to bid using the SMS message system. During the auction roll out, operators were available to take calls from the truckers at all times if they had trouble placing their bid.

[9] The subsidy intervention was conducted on a small subset of the Dakar population and is described and analyzed in **?** and **?**.

Table 1: Summary statistics from the auction platform

(a) Summary statistics all auctions, 2013-2016

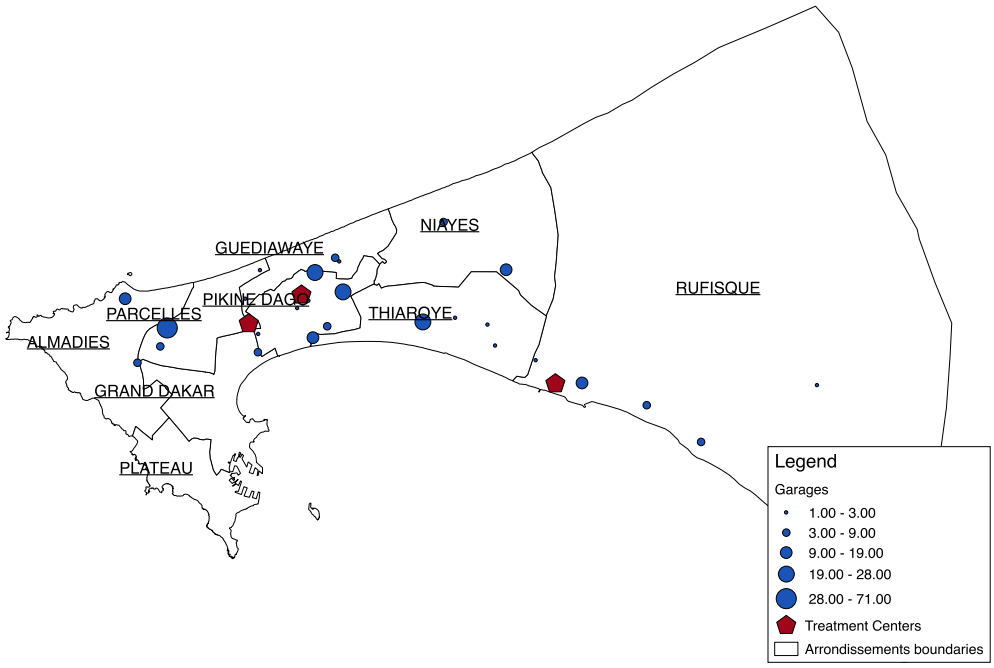
VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Revisable auction indicator	5,162	0.497	0.500	0	1
Number of invited bidders per auction	5,162	12.79	2.672	5	21
New platform indicator (2015-2016)	5,162	0.454	0.498	0	1
Zero bids indicator	5,162	0.191	0.393	0	1
Number of valid bids (at least one bid)	4,177	2.437	1.436	1	11
Number of auction invitations per bidder	126	524.0	354.5	15	1,693
Bidder participation probability	126	0.104	0.136	0	0.569
Number of auction invitations per active bidder	39	886.4	286.6	414	1,693
Active bidder participation probability	39	0.246	0.134	0.0691	0.535

(b) Summary statistics on sealed-bid auctions, 2013-2015

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Number of invited bidders per auction	1,350	14.34	1.848	9	19
Zero bids indicator	1,350	0.0756	0.264	0	1
Number of valid bids (at least one bid)	1,248	2.875	1.475	1	9
Number of auction invitations per bidder	108	179.2	120.9	4	662
Bidder participation probability	108	0.121	0.159	0	0.665
Number of auction invitations per active bidder	39	282.5	135.6	51	662
Active bidder participation probability	39	0.268	0.173	0	0.665

Note: Table uses auction platform data at the auction level, the bidder level, and the active bidder level. Active bidders are those who submit more than 30 valid bids. Table shows random variation in auction format (revisable vs sealed-bid), and the number of invited bidders per auction. It shows an indicator for whether the auction had zero bids, and the number of valid bids conditional on there being at least one. It shows the number of auction invitations and participation probability for all bidders and then again for active bidders.

Figure 1: Distribution of garages and treatment centers in Dakar



trucker is affiliated with the AASS. One caveat with this variable is that the survey participants often provided conflicting answers to some of the key question related to firm organization. Trucks are used by multiple drivers and the arrangements between owners, managers, and drivers are unclear. The different parties seem to have limited information about one another (for example their Association membership or revenue and costs). For example, we asked drivers, managers, and owners about both individual and company-level membership in the Association in the baseline and endline surveys. Over 60% of trucks had some form of conflicted answer about membership. This makes it difficult to use the survey to describe financial arrangements and ownership.

Figure 1 illustrates the location of the main garages and treatment centers in Dakar. Most of our analysis is concentrated in the four main peri-urban Arrondissements Dakar: Pikine, Guediawaye, Niayes, and Thiaroye. We group the four Arrondissements closer to the historical center of the city into one group (Downtown): Plateau, Grand Dakar, Almadies and Parcelles.

The size of each circles illustrates the size of each garage, as measured by the number of affiliated truck. The largest garage (Foire) includes over 70 trucks, while many informal garages include only a handful of trucks. The larger garage tend to be more organized, and controlled by one or a small number of companies affiliated with the main trade association (AASS). The official role of the Association is to help operators collaborate on the procurement of truck parts and to assign large lucrative government and commercial contracts. The influence of the Association extends beyond members and affects the provision of residential contracts by all companies throughout the city. This is because AAAS is involved in the largest garages (where non-member trucks also tend to park), and distributes contracts and services to member and non-member companies. The threat of being excluded represents a risk of reduced profits due to the loss of non-residential contracts and more difficult access to truck parts for

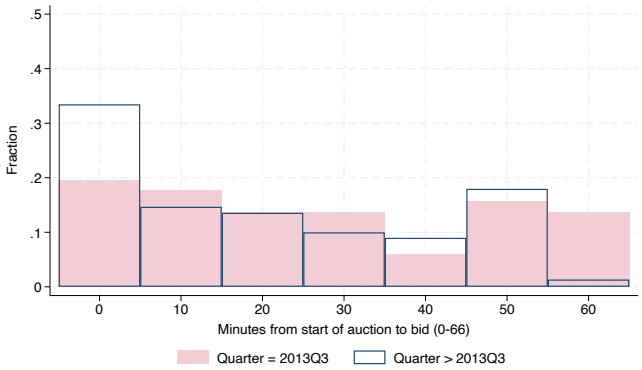
both members and non-members.

1.4 Key stylized facts about the auction platform

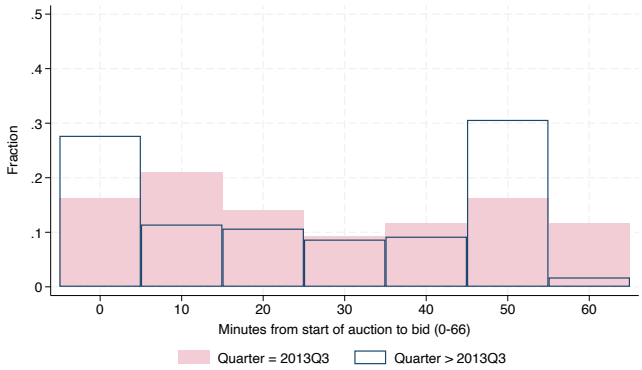
Feature 1: Learning and strategic bidding

Figure 2: Differences in bidding strategy over time and across formats

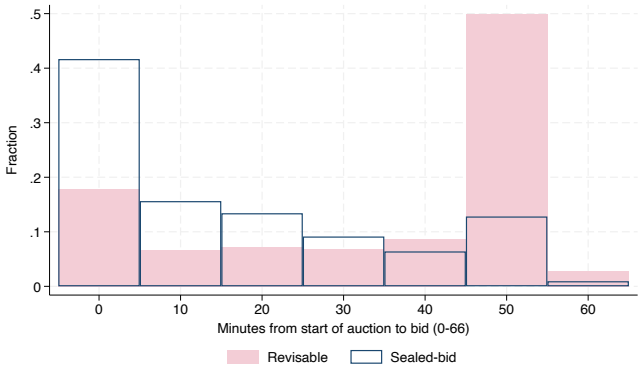
(a) Timing of bids in sealed-bid auctions



(b) Timing of bids in revisable bid auctions



(c) Timing of winning bids across formats after 2013q3



The auction platform represents a new strategic environment for most bidders, and firms likely changed their bidding strategies over time. Figures 2a and 2b provide evidence that this learning process occurred relatively quickly. A salient feature of the platform is the timing of bid. Invited bidders received regular reminders, and, in the case of the revisable auctions, information about the current lowest standing bid. Because of firms' reliance of round numbers and the high probability of ties, firms learned quickly that bidding early in the sealed-bid auction was strategically important. Figure 2a illustrates a stark difference between auctions conducted during the first quarter, and auctions performed in 2014. The fraction of bids received during the first 10 minutes increased from 20% to more than 30%. Although the figure pools together auctions conducted in 2014, the fraction of bids received in the first 10 minutes remained roughly constant after the pilot phase.

In revisable auctions, competitive firms have a strategic advantage of bidding late and undercutting the lowest bid. This is because the revisable auction turns into a sealed-bid auction after minute 50. Figure 2b shows that in first phase of the platform, the distribution of bid times was very similar across the two auction formats. Starting in 2014, bid times converged to bi-modal distribution with more than 30% of bids received immediately after the last message (minute 50). Interestingly, bidders learned that bidding close to the 60 minutes deadline lowered their winning probability, and the mass of bids in the last five minutes of the auction decreased significantly.

These two strategies have implications for the probability of winning. In revisable auctions, the winning bid was received after the last message in over 50% of auctions. In contrast, 40% of sealed-bid auctions cleared in the first 10 minutes. This is consistent with the presence of a group of competitive bidders who quickly learned to bid strategically.

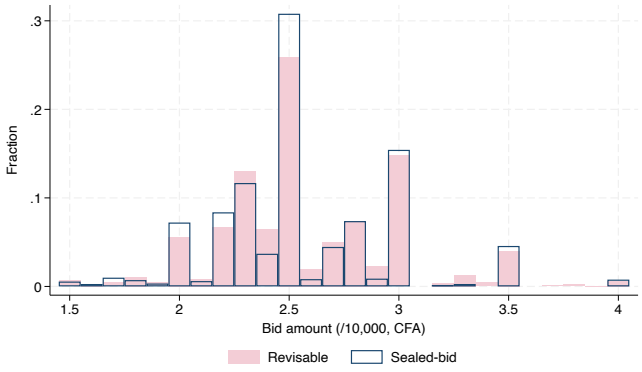
Feature 2: Focal prices and ties

Another important feature of the auctions is the use of focal prices by bidders. As we discussed in **??**, this practice is also common in the regular market, with a modal transaction price of CFA 25,000 across all neighborhoods outside of Rufisque. Figure 3a illustrates this point in the auctions by plotting the distribution of accepted and rejected offers in the sealed-bid auctions. Nearly 50% of winning bids in the sealed-bid auctions are equal to 25,000 or 30,000. This feature is present in the revisable auctions also to a lesser extent in the revisable auctions. The ability to undercut the current lowest bid allows competitive bidders to bid 24,000 or 23,000 at a higher rate. In contrast, competitive bidders in the sealed-bid auctions are more likely to submit lower bids (e.g. 20 or 22). As a result the average winning bid in the two formats is nearly identical (24,559 vs 24,621).

The over-reliance on focal prices leads to a high frequency of ties. Figure 3b plots the evolution over time of the probability of observing a tie-winning bid across both formats bid (conditional on receiving at least two bids). After the pilot phase, the tie probability is 25% in sealed-bid auctions and 15% in revisable auctions.

Figure 3: Distribution of winning bids and ties

(a) Winning bids across formats



(b) Tie probability

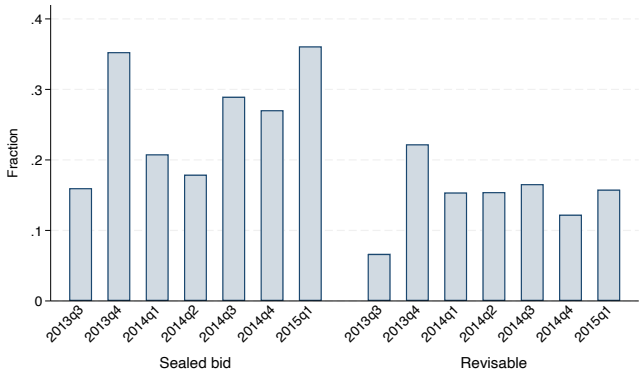
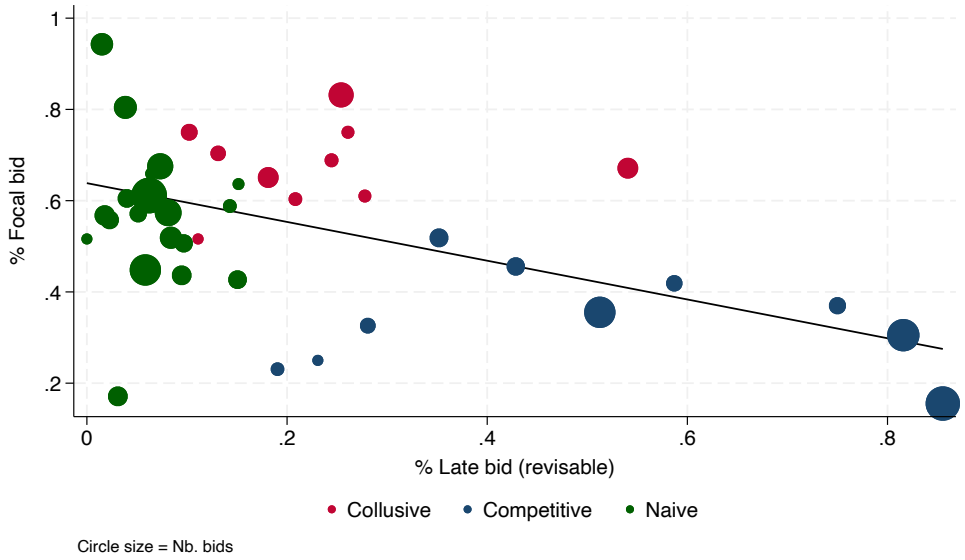


Figure 4: Bidder competitive conduct classification



Feature 3: Heterogeneity in competitive conduct

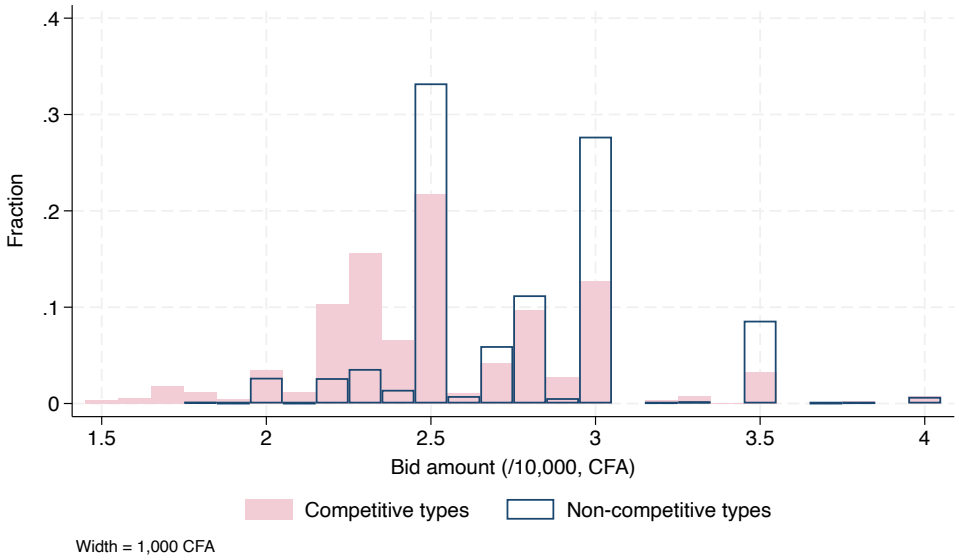
While the analysis of bid times is consistent with strategic bidding and profit maximization from some bidders, the fact that a significant fraction of bidders place early bids in the revisable auction suggests that many bidders do not use profit-maximizing strategies. In Houde et al. (2024) we argue that this is a reflection of the prevalence of tacit collusion in the market. In particular, we use a k-mean clustering algorithm to classify bidders between one competitive and two non-competitive groups based on their propensity to avoid dominated strategies on the platform.

Figure 4 illustrates the result of this classification by plotting the probability of submitting a focal bid in sealed-bid auctions (y-axis) against the propensity of bidders to submit a late bid in revisable auctions (x-axis). We measure both variables for the 39 active bidders. The figure illustrates a clear negative correlation between the use of focal prices and the probability of snipping (the size-weighted correlation coefficient is -0.59). Competitive bidders (blue) are more likely to snipe and under-cut the lowest standing bid and avoid the two focal prices (25K and 30K) in the sealed-bid auctions. The remaining group of bidders avoid competition by bidding early and focusing mostly on focal prices. Both strategies soften price competition on the platform and lead to higher offers to consumers and a lower take-up rate.

Figure 5 illustrates the bidding strategy of competitive and non-competitive types in the sealed-bid auction format. We focus on the sample auctions conducted before the platform re-design, and after the pilot phase. Relative to non-competitive types, competitive bidders submit lower bids by more frequently bidding 22,000 and 23,000, and rely on a richer price grid. Both features imply that competitive bidders are more likely to win, and less likely to tie.

Despite this higher level of sophistication by competitive types, it is clear that all bidders rely on a finite number of bids, which inconsistent standard models of auctions. In this paper we remain somewhat agnostic about what causes the presence of mass-points, and assume that bidders maximize their

Figure 5: Distribution of bids in the sealed-bid auctions by bidder types



expected profits by selecting a bidding strategy from a coarse pricing grid. To account for differences in competitive conduct across bidders, we construct a bidder-specific consideration-set $\mathcal{B}_i$ that is common across all auctions.

This assumption is compatible with models of tacit collusion, as well as the existence of pricing frictions and/or bounded rationality.[10] Under the collusion interpretation, since bidders are randomly invited to participate, they have limited opportunities to select a "winner" the ring prior to bidding (as in Asker (2010) for instance). Instead, collusion most likely takes the form a tacit agreement either not to participate at all, or submit common focal prices in order to prevent undercutting, and target the expected reservation price of consumers. See **?** for a discussion of the tacit collusion interpretation.

We construct the consideration set of each bidders by identifying their most common bids. In particular, we define the a "base" choice-set as the five most common prices used by all bidders: 22K, 25K, 28K, 30K, and 35K. For each bidder, we then augment this set with any bids that are chosen with a frequency greater than 5% of the maximum frequency for this specific bidder.[11] The advantage of this approach is that bidders are differentiated in their levels of sophistication and competitive conduct, which allows us to rationalize the bidding strategy of all bidders.

Table 2 summarizes the consideration set of competitive and non-competitive bidders. As expected, competitive bidders have richer consider sets (15.80 distinct price points on average), and use a bidding strategy that are less concentrated on a handful of grid points (as measured by the Choice HHI variable). Importantly, competitive bidders are not necessarily more active on the platform. While they tend to

[10] One explanation is provided by Beaman, Magruder, and Robinson (2014), who show that most people do not have the kinds of denominations necessary to conduct trade at very precise prices.

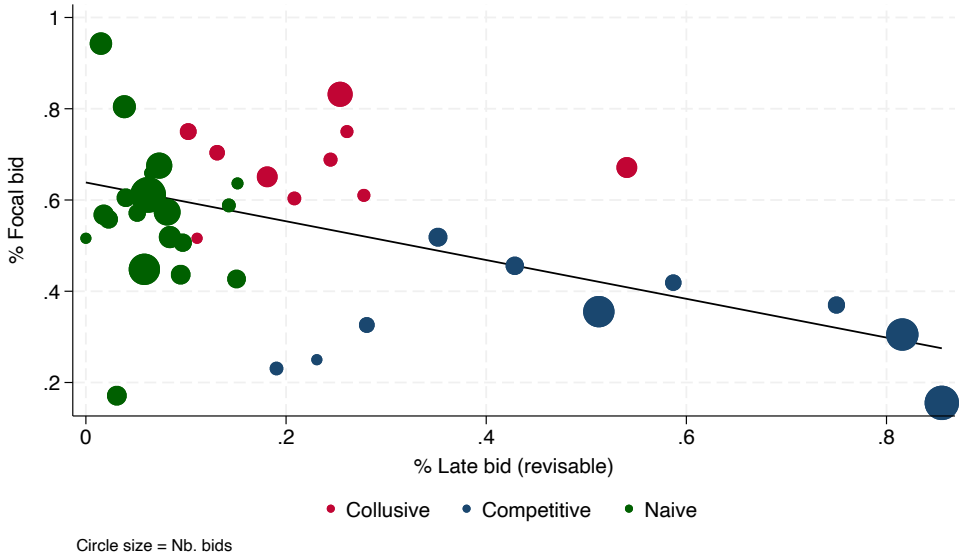
[11] We use the ratio of the frequency over the max frequency instead of frequency level because some of the active bidders have near uniform grid (i.e. almost continuous). For those bidders, the frequency is going to be uniformly low, and selecting the consideration set based on a minimum frequency criteria would drop bids that are considered "relatively" frequently for that bidder.

Table 2: Summary statistics of bidder's consideration sets by bidder types

	Competitive		*Naive*		*Collusive*	
	Mean	Std.Dev.	Mean	Std.Dev.	Mean	Std.Dev.
Choice HHI	0.18	0.06	0.26	0.07	0.34	0.06
Nb. Options	15.80	8.38	10.16	2.50	8.00	2.71
Avg. Bids	2.63	0.262	2.72	0.18	2.89	0.20
Nb. Bids	136.20	121.07	122.95	82.83	90.80	55.77
Nb. Bidders	10.00	0.00	19.00	0.00	10.00	0.00

Note: Bids are measures in 10,000 CFA.

Figure 6: Distribution of bidding strategies across provider types



submit bids in more auctions (135 on average), bidders in the other two groups are also frequent bidders (122.95 and 90.80 respectively for naïve and collusive types).

2 Empirical model of bidding in the sealed-bid auction

In this section, we provide a standard model of procurement auctions that coincides with the platform deployed in Senegal, and then describes an identification and estimation strategy.

2.1 Model description

There are $i = 1, ..., I$ desludgers competing in auctions $t = 1, ..., T$. Since each auction is associated with a single client, we use t interchangeably to index auctions and clients. During each auction t, a random subset A_t of desludgers are invited to bid. Bidders know the number of invited bidders, but not their

identity. The timing of auction t is (1) the platform invites a subset A_t of desludgers to enter, (2) the desludgers simultaneously and independently decide whether to enter or not, (3) those desludgers who enter submit competitive bids, (4) the lowest bid is presented to the household that initiated the auction, who accepts or rejects; ties are broken randomly. The auction is paid-as-bid, so a winning desludger is paid his bid by the household.

When faced with the winning bid, b_t^*, the household accepts if its willingness-to-pay w_t is larger than the log price, or

$$w_t \geq b_t^*. \tag{1}$$

This is equivalent to a paid-as-bid procurement auction with a secret reserve price. Since the household is simply presented a take-it-or-leave-it offer at the end of the process, she has no reason to deviate from this decision rule, and we model the neighborhood-household as a demand curve generating stochastic draws w_t and not a strategic player. Let the demand curve is given by the expected take-up probability:

$$D(b|x_t) = Pr[w_t \geq b].$$

where x_t is a vector of characteristics of client t that affects the distribution of willingness-to-pay (including for instance the neighborhood, the period, etc.).

Each desludger i who enters auction t draws a cost γ_{it} which perturbs his mean cost for completing job t, $\bar{c}_{it}$, giving him a total cost of $c_{it} = \bar{c}_{it} + \gamma_{it}$. This private-value shock is independently distributed across bidders and auctions according to a continuous density function $f(\cdot)$ that is common across bidders.

We assume that each bidder i selects bids in consideration set $\mathcal{B}_i$ to maximize expected profits:

$$\pi_i(c_{it}|x_t, n_t) = \max_{b_{it} \in \mathcal{B}_i} G_{it}(b_{it}|x_t, n_t)(b_{it} - c_{it}), \tag{2}$$

where $G_{it}(b_{it}|x_t, n_t)$ is the probability that b_{it} is selected as the winning bid and the household accepts b_{it}, or

$$G_{it}(b_{it}|x_t) = E\left[1\{b_{it} = B_{-i,t}\} \times \frac{1}{1 + K_t(b_{it})}\Big| x_t, n_t\right] \cdot D(b_{it}|x_t). \tag{3}$$

where $B_{-i,t}$ is lowest-bid among rival invited bidders, and $K(b_{it}) = |\{i \in A_t : b_{it} = \min_{k \in A_t \setminus i} B_{kt}\}|$ is the number of rival bids tied with b_{it} in the auction. These probabilities are determined in equilibrium and depend on the underlying distributions of bidder characteristics. Note that the belief function differs across bidders due to observed differences in cost (e.g. distance), as well as differences in consideration sets.

To decide whether to enter auction t, each desludger compares his expected profits with his stochastic opportunity cost:

$$\bar{\pi}_{it} = \mathbb{E}_{c_{it}}\left[\pi_i(c_{it})|x_t, n_t\right] > \kappa_{it} + \sigma_\eta \eta_{it}. \tag{4}$$

If the desludger decides to enter the auction, he then learns the specific cost he faces that day and decides what bid to place. If the desludger submits a bid, he cannot take on other work until the auction is concluded, which means the outside option is a "sunk cost" forfeited at the time he determines the cost of the job and prepares his bid. The shock is allowed to depend on characteristics of the desludger

and clients. If $H(\cdot|z_{it})$ denotes the CDF of η_{it} conditional on observed bidder and client characteristics z_{it}, the equilibrium probability of participation is given by:

$$\Pr(\sigma_\eta \eta_{it} < \bar{\pi}_{it} - \kappa_{it}|z_{it}) = H(\bar{\pi}_{it} - \kappa_{it}|z_{it}) \tag{5}$$

The opportunity cost $\kappa_{it} + \sigma_\eta \eta_{it}$ need not be positive: a desludger can lose money by sitting idle with valuable capital, paying wages to workers and not doing any work.

Finally, we assume that the shocks w_{it}, η_{it} and γ_{it} are to be independent distributed conditional on x_t.

2.2 Identification and estimation strategy

This section discusses the identification and estimation the key pieces of the model: the distribution of willingness-to-pay, w_{it}; the average observable costs, $\{\bar{c}_{it}\}_{i,n,t}$; the outside options κ_i; the distribution of cost shocks, $F(\gamma_{it})$; and the distribution of entry shocks $H(\eta_{it})$ where firms (bidders) are indexed by i, and clients (auctions) are index by t. We estimate the model in three stages, starting with the demand function.

Demand estimation

We parametrize the distribution of willigness-to-pay using a normal distribution, which leads to a binary choice model of demand:

$$D_t(b_t^*|x_t) = \Pr(\varepsilon_t > \alpha b_t^* - x_t\beta) = 1 - \Phi\left(\alpha b_t^* - x_t\beta\right) \tag{6}$$

where $\varepsilon_{nt} \sim N(0,1)$, and $\alpha = 1/\sigma_\varepsilon$ measures heterogeneity in taste across clients. Since the lowest bid, b_t^*, is determined in the market, there is concern that it might be endogenous, and direct estimation of (6) by maximum likelihood will result in biased estimates. However, any variable that significantly shifts the distribution of clearing prices and is not revealed to the household is a valid instrument. In particular, the characteristics of the auctions themselves that are unobserved by the household or directly by bidders, but do influence the distribution of clearing prices provide a variety of potential instruments.

To address this potential simultaneity problem, we exploit random variation from invitation lists to estimate the following system of equations:

$$b_t^* = z_t\pi + \sigma_\eta u_t \tag{7}$$

$$d_t = \mathbb{I}\{x_t\beta - \alpha b_t^* + \varepsilon_t \geq 0\}, \tag{8}$$

where (u_t, ε_t) are distributed according to a joint normal distribution with correlation coefficient ρ. The variable z_t includes the covariates x_t and a set of instrumental variables, which affect the second stage only through the realized clearing price, b_t^*. We can then test the hypothesis of exogenous prices by testing the null hypothesis that ε_t is independent of u_t (or ρ_t). We construct two instruments using variation from the random-invitation lists: (i) number of active bidders invited to auction t, and (ii) the

minimum distance among active invited bidders to the client.

Estimation of the variable cost distribution

Because of the presence of ties, the standard first-order condition approach for estimating firms' costs from bid data would lead to a rejection of the model, since firms using focal bids can strictly increase their profits by lowering their bid by a small amount. We adopt a slightly different approach to cost estimation that draws on the literature on revealed preference inequalities.

Since $D_t(b)$ is monotonically decreasing in b, the probability of a transaction is monotonically decreasing in the bid amount: $G_i(b'|x_t, n_t) > G_i(b|x_t, n_t)$ for all $b' < b$. This does not guarantee however that the expected profit is strictly concave in b. In particular, because of the presence of mass points in the bid distribution, some bids can be strictly dominated options. To simplify the notation, we omit the dependence of the winning probability with respect to (x_t, n_t): $G_i(b|x_t, n_t) \equiv G_{it}(b)$.

If b_i chosen it must satisfy the following two conditions:

$$\text{For all } b' > b_i: \qquad G_{it}(b_i)(b_i - c) > G_{it}(b')(b' - c) \rightarrow c < \frac{G_{it}(b_i)b_i - G_{it}(b')b'}{G_{it}(b_i) - G_{it}(b')} = \mu_{it}(b_i, b') \tag{9}$$

$$\text{For all } b' < b_i: \qquad G_{it}(b_i)(b_i - c) > G_{it}(b')(b' - c) \rightarrow c > \frac{G_{it}(b')b' - G_{it}(b_i)b_i}{G_{it}(b') - G_{it}(b_i)} = \mu_{it}(b', b_i) \tag{10}$$

The ratio $\mu_i(b', b_i)$ measures the "marginal revenue" associated with choosing b_i (high) over b' (low). When firms reject high-bid alternatives it is because revenue gain from that alternative outweighs the lower transaction probability overall. This yields to an upper bound on the cost. The lower bound is obtained by looking at rejected bids that are lower than the chosen option. The firm chose to place a higher bid in this case because to increase the revenue per transaction. Note that a bid b_i cannot be rationalized by any cost if the two bounds cross. Once that happens, the chosen bid b_i is a dominated option and is not rationalizable. This condition can be verified empirically using the empirical analogue of the winning probability function. In practice, we find that 6% of bids are not rationalizable by consideration sets $\mathcal{B}_i$. We estimate the model by maximum-likelihood focussing on the sample bids that satisfy this condition.

Recall that the ordered consideration-set for bidder i is given by: $\mathcal{B}_i = \{b_1, b_2, \ldots, b_J\}$ where $b_1 < b_2 < \cdots < b_J$. Let $Y_{it} = \{y_{it,1}, y_{it,2}, \ldots, y_{it,J}\}$ denotes a vector of choice indicator variables, such that $y_{it,j} = 1$ if $b_i = b_{ij}$, and $\sum_j y_{it,j} = 1$. Similarly, let $\mu_{it,j} = \mu_{it}(b_{j-1}, b_j)$ denotes the marginal revenue of choosing bid $j-1$ over alternative j in the set.

For bids that are rationalizable, the optimality conditions for interior bids can be written as:

$$y_{it,j} = 1 \text{ if } \rightarrow \mu_{it,j-1} = \frac{G_{it}(b_{j-1})b_{j-1} - G_{it}(b_j)b_j}{G_{it}(b_{j-1}) - G_{it}(b_j)} \leq c \leq \frac{G_{it}(b_j)b_j - G_{it}(b_{j+1})b_{j+1}}{G_{it}(b_{j+1}) - G_{it}(b_{j+1})} = \mu_{it,j} \tag{11}$$

For bids in at the edge of the consideration set the upper and lower bounds are defined either by zero

or the chosen bid:

$$y_{it,0} = 1 \text{ if } \rightarrow \mu_0 = 0 \leq c \leq \frac{G_{it}(b_1)b_1 - G_{it}(b_2)b_2}{G_{it}(b_1) - G_{it}(b_2)} = \mu_1 \tag{12}$$
$$y_{it,J} = 1 \text{ if } \rightarrow \mu_{J-1} = \frac{G_{it}(b_{J-1})b_{J-1} - G_{it}(b_J)b_J}{G_{it}(b_{J-1}) - G_{it}(b_J)} \leq c \leq b_J = \mu_J$$

Since the cost cannot be inverted from the chosen bids we use a semi-parametric functional from to estimate the cost distribution. In particular, we use a single index function for c_{it}:

$$c_{it} = x_{it}\beta + \gamma_{it} = \bar{c}_{it} + \gamma_{it} \tag{13}$$

where x_{it} is a vector of observed cost shifters (to the econometrician), and γ_{it} is an privately observed idiosyncratic cost residual of performing the contract distributed according to cumulative distribution function $F(\gamma)$. Due to the single index property of the cost function, the model corresponds to an ordered discrete choice model with non-parametric distribution of the residual (as in Klein (1993)). We use a mixture of distribution (MOD) estimator to approximate $F(\cdot)$ (Coppejans, 2001).

Let θ denote the vector parameters determining the distribution function $F(\gamma|\theta)$. We estimate the model using maximum likelihood:

$$\min_{\beta,\theta} \quad \sum_i \sum_j \sum_{j=1}^{J_i} y_{it,k} \log\left[F(\mu_{it,j} - x_{it}\beta|\theta) - F(\mu_{it,j-1} - x_{it}\beta|\theta)\right] \tag{14}$$
$$s.t. \quad F(\gamma|\theta) = \sum_{j=1}^{h} \pi_j \Phi((\gamma - \eta_j)/\sigma_j)$$

where $Y_{it,j}$ is an indicator variable equal to one if firm i submit bid b_k in auction j (zero otherwise), J_i is the number of choices available to bidder i, $\theta = \{\pi_j, \eta_j, \sigma_j\}_{j=1,\ldots,h}$ is the vector mixing distributional parameters, and $\Phi(\cdot)$ is the normal CDF.

Unlike in standard discrete-choice models, we do not need to impose a scale restriction on γ since the cutoffs points (i.e. μ_k) are measured the local currency, and are calculated outside of the bid choice model using the empirical distribution of winning probabilities. To identify a non-parametric distribution of γ we also need that the support of the index, $\mu_k - x_{it}\beta$, is at least as large as the support of γ. This is achieved if at least one x_{it} varies continuously across bidders and auctions, or if the observed expected cutoffs exhibit sufficient variation across auctions. Both conditions are satisfied in our context because x_{it} includes distance to the client, and the inclusion of neighborhood and time fixed-effects in the demand equation insures that the cutoffs vary substantially across auctions. In practice, we do not impose an upper and lower bonds on the support of γ, which leads to the possibility of negative cost realizations. Although we do not impose a bounded support restriction in the estimation, we use a truncated distribution when computing the counter-factual auction outcomes.

To estimate bidders' beliefs about the probability of winning the auction (and therefore the marginal revenue cutoffs), we follow the approach pioneered by Guerre, Perrigne, and Vuong (2000). The information set x_t summarizes what is common knowledge by all bidders invited to participate in the auction. This includes the number of potential bidders, the location of the client, the auction date, and

the auction timing. The auction date is informative both because of seasonalities in the demand and supply, and because of potential learning effects which affect the evolution of beliefs. The auction time is a strong predictor of the number of bids because it affects the capacity utilization of firms and whether or not the service can be fulfilled during the same day.

To account for the presence of mass points in the bid distribution, we use a resampling approach to calculate $G(b|x_t, n_t)$ (similar to Hortaçsu and McAdams (2010) and Hortaçsu and Kastl (2012)). In particular, we resample the outcome of auctions with similar information sets, then calculate the empirical frequency of winning the auction with a given bid. Consider an auction t with characteristics: $x_t = (\text{lat}_t, \text{long}_t, \text{Hour}_t, \text{Date}_t, N_t)$. The probability sampling auction t' is a function of the distance in characteristics, measured using a normal density kernel:

$$K(x_t, x_{t'}) = K(d_{t,t'}/h_d)K((\text{Hour}_t - \text{Hour}_{t'})/h_{hour})K((\text{Date}_t - \text{Date}_{t'})/h_{Date})K((N_t - N_{t'})/h_N)$$

$$\rightarrow \omega_{t,t'} = \frac{K(x_t, x_{t'})}{\sum_{t'} K(x_t, x_{t'})}$$

The bandwidth parameters are set equal to twice the "Silverman" rule of thumb for the optimal bandwidth for univariate kernels. This is somewhat arbitrary, but it ensures that the same auctions are not repeatedly sampled (i.e. over-smoothing). Some auctions that take place in low density neighborhoods end up with nearly uniform weights.

Our main specification allows bidders to have heterogeneous beliefs about the probability of winning. We do this by removing bidder i's bids from the "bid-to-beat" distribution. We proceed as follows:

i. Exclude all bids from bidder i from the data-set

ii. Assign a sampling probability to other bids equal to $\omega_{t,t'} \times 1/n_{t'}$ (i.e. assume that rival bids are sample uniformly conditional of drawing bids from auction t'). Let $\omega_{t,t',j}$ denotes the probability of sampling bid j.

iii. Sample a number of invited bidders n consistent with the distribution of potential bidders

iv. Sample $n-1$ observations from the multinomial distribution $\omega_{t,t',j}$.

v. Rank bids and compute the number of ties

vi. For each value of bids b in the grid, calculate the outcome of the auction: lose, tie or win.

vii. Repeat steps 3-6 10,000 times, and average the outcomes to calculate the frequency at which the bidder wins each auction.

These steps evaluate the probability of winning for a bidder participating in auction t. This process is repeated for all bidders and auctions.

The second approach pools data across bidders, and assumes that bidders have symmetric beliefs. We use the empirical distribution of the second-lowest bid to evaluate the probability of winning the auction. This is done using the following algorithm:

i. Sample an auction t' from multinomial distribution $\omega_{t,t'}$

ii. Extract the second-lowest bid in that auction (ignoring the possibility of more than one tie).

iii. For each value of bids b in the grid, calculate the outcome of the auction: loose, tie or win.

iv. Repeat steps 3-6 10,000 times, and average the outcomes to calculate the frequency at which the bidder wins each auction.

Estimation of the participation cost distribution

The decision to submit an offer or not is based on the expected profit from the auction and the value of the next-best alternative for firm i. Let $\kappa_{it} = z_{it}\delta + \eta_{it}, \eta_{it} \sim N(0, \sigma_\eta^2)$ denote the outside option of firm i when invited to bid for client j. Recall that we assume that firms decide to submit an offer before observing the idiosyncratic component of their variable profits, γ_{it}. If $\beta_{it}(c)$ denotes the optimal bidding strategy summarized in equation (11) and (12), the expected profit associated with submitting an offer is given by:

$$\begin{aligned} \bar{\pi}_{it}(\bar{c}_{it}, \mathcal{B}_i) &= \int G_{it}(\beta_{it}(c))\,(\beta_{it}(c) - c)\, f(c - \bar{c}_{it}|\theta)dc \\ &\equiv \int \left\{ \max_{b \in \mathcal{B}_i} G_{it}(b)\,(b - c) \right\} f(c - \bar{c}_{it}|\theta)dc \end{aligned} \tag{15}$$

In practice, the expected profit is approximated using quadrature. For each value of c over the integration grid, the profit is replaced by the maximum profit among the bids in each firm's consideration set.

The participation decision is therefore given by the following discrete-choice problem:

$$D_{it} = \begin{cases} 1 & \text{If } \bar{\pi}_{it} - \kappa_{it} > 0 \\ 0 & \text{Else.} \end{cases} \tag{16}$$

This yields to a probit model with the following log-likelihood function:

$$\max_{\tilde{\delta}} \sum_t \sum_{i \in A_t} D_{it} \log\left(1 - \Phi(-w_{it}\tilde{\delta})\right) + (1 - D_{it}) \log\left(\Phi(-W_{it}\tilde{\delta})\right) \tag{17}$$

where $w_{it} = (\bar{\pi}_{it}, -z_{it})$, and $\tilde{\delta} = (1/\sigma_\eta, \delta/\sigma_\eta)$.

3 Estimation results

3.1 Demand and willigness-to-pay

We estimate the demand function using the full sample of auctions for non-subsidized clients (roughly 500). That is we combine auctions before and after the design change, as well as auctions conducted under the sealed-bid and revisable format. These changes are orthogonal to clients' willingness-to-pay, and generate rich variation in prices.

Table 3 presents the estimation results for the demand function. To test for the endogeneity of prices, we compare the estimation results obtained using a probit and linear-probability model, with results

Table 3: Estimation results for the demand function under alternative identification assumptions

VARIABLES	(1) OLS	(2) IV	(3) Probit	(4) IV-Probit	(5) IV-Probit	(6) IV-Probit
Winning bid (x1000)	-0.29[a]	-0.38[a]	-1.07[a]	-1.31[a]		
	(0.019)	(0.13)	(0.080)	(0.43)		
Active bidders (<30) x Less 10 km					-0.022[a]	
					(0.0038)	
Min. distance active (<30)					0.011[a]	
					(0.0024)	
Observations	3,620	3,620	3,620	3,620	3,620	3,620
R-squared	0.101	0.096				
Mean dep. variable	0.30	0.30	0.30	0.30	0.30	0.30
F-test (invitations)		44.3		44.3	44.3	44.3
Over-identification test (chi2)		0.99				
Correlation coefficient				0.083	0.083	0.083
Exogneity test (p-value)				0.58	0.58	0.58
Elasticity			-3.52			
Average WTP			2.01			

Robust standard errors in parentheses
[a] p<0.01, [b] p<0.05, [c] p<0.1

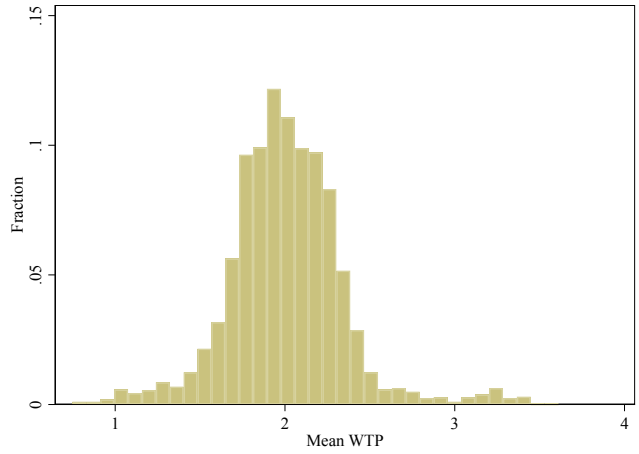
obtained using our invitation-list instruments. Recall that we use two instruments: (i) the number of active bidders located within 10 kilometers of the client, and (ii) the minimum distance among active bidders. Since active bidders are more likely to bid and tend to bid more competitively, these two variables are strong predictors of the winning bid. We use data from both auction formats to estimate the demand function since consumers are randomly assigned.

Column (5) presents the results of the first-stage regression. The two variables are strongly correlated with the offer that consumers receive (F-test=43). The sign of the coefficients confirms that inviting more and close-by active bidders lower prices. Columns (2) and (4) show that we cannot reject the null hypothesis of exogenous prices. The chi-square statistic associated with the over-identification test obtained from the linear-probability model confirms that the invitation variables are orthogonal to the residual determining the WTP of consumers. The exogeneity-test from the IV-Probit specification estimates that the correlation between the price and WTP residual is indistinguishable from zero (p-value=0.89). See equations (7) and (8). This suggests that the unobserved factors determining the decision of bidders to participate in the auction and the value of the bid are uncorrelated with the Willingness to Pay of consumers. This is not unreasonable in our context since the randomization in invitations and auction formats generates the bulk of the variation in price, and consumers do not decide to call the platform knowing this information.

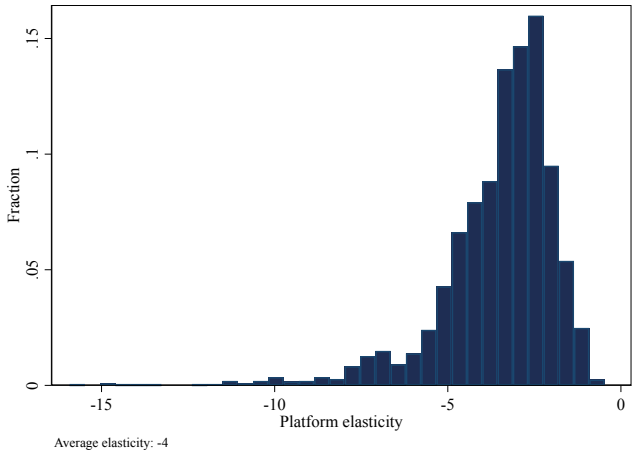
Figure 7 further explore the heterogeneity in demand across consumers. Figure 7a plots the distribution of the mean WTP among consumers who accepted an offer (i.e. $E[w_t|x_t]$). This distribution is highly concentrated between 20,000 CFA and 25,000, suggesting that the consumer surplus generated

Figure 7: Distribution of WTP and demand elasticity in the platform

(a) Distribution of WTP in the platform



(b) Distribution of consumer elasticity in the platform



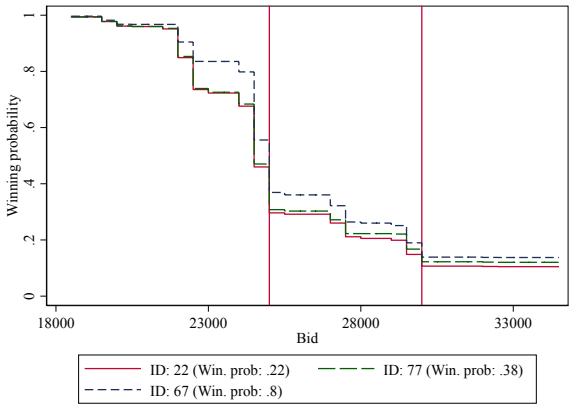
by the platform varies between 1,000 and 3,000 for most consumers.

The demand for desludging on the platform is very elastic (-3.34 on average), but also highly heterogeneous. A significant portion of the demand elasticities are lower than -5. There are two main reasons for this high observed elasticity of demand among households. First, the mechanical service is expensive and there is a cheaper substitute available in manual desludging. Second, if a given household fails to purchase from the platform, they can purchase a desludging in the market from another provider; they may also even use the price quoted on the platform to negotiate. Both of these channels contribute to the challenges faced by the platform in pleasing and retaining customers and are a key problem that the platform must consider. The results suggest that the outside option of most consumers who call the platform is generated mostly from the distribution of negotiated prices in the regular market, which is equal to 25,000 on average.

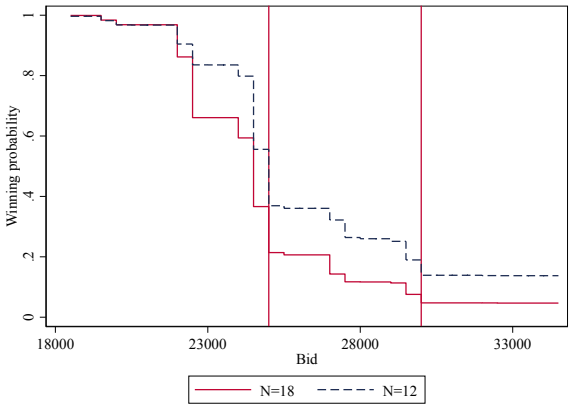
Table 4 summarizes the distribution of willingness-to-pay and elasticity across neighborhoods, as well as the average bids and acceptance probability. The last column measures the average transaction prices in the market obtained from a large survey of household desludging choices (see **?** for a discussion of this data-set). Across most neighborhoods the average offers (b_{min}) are larger than the average negotiated prices (2.57 vs 2.51). On the other hand average accepted bids are on average CFA 700 lower than the average transaction prices (or 2.8%). There are two exceptions to this. Callers from Rufisque have a very favorable outside option, because of the highest level of competition in this Arrondissement, and as a result accept a only 13% of offers (for an average transaction price of $16,700). Clients from the Niayes neighborhood receive offers that are on average 3,500 higher than the market average, and as a result accept only 21% of platform offers. This low take-up is due to the fact that Niayes is located further away from most providers active on the platform.

Figure 8: Estimated belief functions

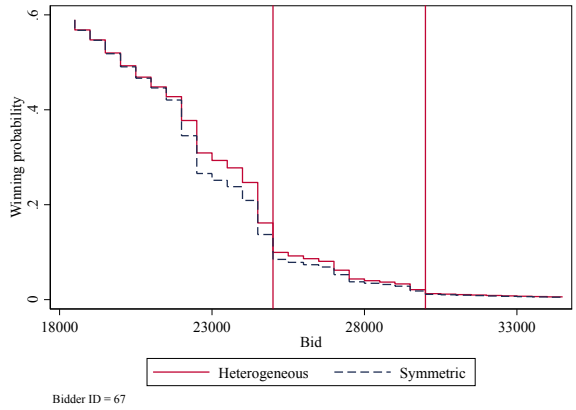
(a) Winning probability across three bidders



(b) Winning probability for high and low number of invited bidders



(c) Transaction probability for two models of beliefs



(d) Transaction probability high and low WTP consumers

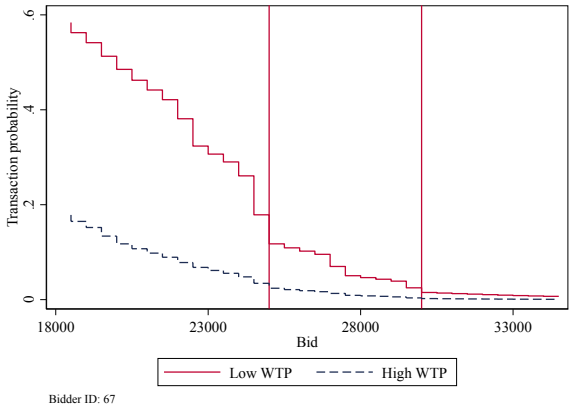


Table 4: Summary statistics on willingness-to-pay and demand elasticities by neighborhoods

Neighborhoods	Avg. auction outcomes						Avg. prices Survey
	N	WTP	Elasticity	$\Pr(d=1\|b_{min})$	$b_{\min}$	$b_{min}\|d=1$	
Downtown*	465	2.32	-2.59	0.44	2.49	2.38	2.78
Guediawaye	549	1.95	-3.24	0.30	2.48	2.43	2.50
Niayes	886	2.03	-4.49	0.21	2.85	2.73	2.50
Pikine	381	1.75	-2.94	0.31	2.25	2.11	2.20
Rufisque	123	1.18	-6.09	0.13	2.59	1.67	1.50
Thiaroye	1217	2.06	-3.22	0.32	2.54	2.46	2.50
Total	3621	2.01	-3.52	0.30	2.57	2.44	2.51

Downtown arrondissements: Almadies, Plateau, Grand Dakar, Parcelles.
Price units: CFA/10,000.

3.2 Cost function results

We estimate the distribution of variable cost by maximum-likelihood (see equation (14)). We impose three additional sample restrictions relative to sample used to estimate of the demand function. First, we keep auctions conducted during the stable phase of the platform to limit mis-specifications of the bidding choice model. In particular, we exclude auctions conducted in the pilot phase (third quarter of 2013), as well auctions conducted after the platform re-design. We drop the pilot phase to limit the importance of learning among invited bidders. Similarly, the changes implemented significantly reduced competition on the platform, which led to a spike in the number of failed auctions. Second, we focus on bids submitted in sealed-bid auctions by active bidders. Our final estimation sample includes 2,737 bids submitted by 39 firms. This corresponds to 10,160 invitations (or 26% participation rate), and 1,302 distinct auctions.

A key input in the estimation of the model is the empirical distribution of winning probability. Figures 8 plots this estimated belief functions used to construct the likelihood function for a subset of auctions and bidders. The top two figures plot the empirical probability of winning the auctions, and the bottom two plot the transaction probability over a 500 grid of bids. The first takeaway from these figures is the importance of accounting for ties and focal prices. The probability of winning the auction decreases very abruptly at 25,000 and 22,000, and to a lesser extent at 30,000. The second takeaway is the importance of accounting for firm asymmetries. As we can see from Figure 8a, Bidder 68 is the most sophisticated bidder (finer consideration set), and also submit on average lower bids. As a result this bidder faces a higher probability of winning at all bid levels. Figure 8c illustrates this point by comparing the transaction probability assuming symmetric (blue) and symmetric beliefs. Finally, Figures 8b and 8d highlight the importance of controlling for observed differences in willingess-to-pay and competition across auctions.

Table 5: Maximum likelihood estimates of the variable and participation cost functions

	(1)		(2)		(3)		(4)	
	Est.	SE	Est.	SE	Est.	SE	Est.	SE
Variable cost:								
Total driving distance (km)	0.019	0.002	0.025	0.002	0.034	0.004	0.025	0.002
1(Association member)	-0.171	0.016	-0.012	0.017	-0.038	0.069	-0.014	0.017
1(Single truck bidder)	0.175	0.017	0.112	0.016	0.250	0.035	0.116	0.016
1(Lunch)	-0.015	0.020	-0.015	0.019	0.050	0.044	-0.016	0.019
1(Afternoon)	0.002	0.027	0.000	0.026	0.021	0.058	0.002	0.026
1(Foire garage)	-0.288	0.026	-0.306	0.024	-0.480	0.054	-0.344	0.033
Client longitude	0.416	0.463	-0.257	0.457	0.814	0.914	-0.270	0.452
Client latitude	3.092	0.820	3.006	0.844	1.685	2.019	3.075	0.812
N. Invited bidders							0.013	0.004
% Same garage							0.120	0.068
FE cluster: collusive			0.359	0.025				
FE cluster: naive			0.297	0.020				
$\ln(1-\pi_1)/\pi_1$	1.269	0.429	-1.242	0.319	-0.406	0.595	1.289	0.326
μ_1	0.176	0.073	-0.140	0.050	0.034	0.057	0.137	0.051
$\ln \sigma_1$	-0.569	0.074	-1.334	0.049	-0.747	0.088	-0.572	0.072
$\ln \sigma_2$	-1.249	0.059	-0.580	0.070	-1.653	0.253	-1.330	0.049
Arrond FE	Y		Y		Y		Y	
Day FE	Y		Y		Y		Y	
Quarter-Year FE	Y		Y		Y		Y	
Bidder FE	N		N		N		N	
Participation probability:								
Expected profits	10.900	0.716	11.500	0.811	9.146	0.853	11.480	0.808
1(Association member)	-0.468	0.034	-0.336	0.032	-0.233	0.077	-0.339	0.032
1(Single truck bidder)	0.286	0.034	0.204	0.033	-0.035	0.071	0.209	0.033
1(Lunch)	-0.065	0.038	-0.064	0.038	0.047	0.076	-0.065	0.038
1(Afternoon)	-0.150	0.051	-0.149	0.051	-0.021	0.103	-0.146	0.051
1(Foire garage)	0.004	0.033	0.026	0.033	0.149	0.068	-0.010	0.057
N. Invited bidders							0.018	0.009
% Same garage							0.119	0.148
FE cluster: collusive			0.359	0.025			0.358	0.025
FE cluster: naive			0.297	0.020			0.297	0.019
Day FE	Y		Y		Y		Y	
Quarter-Year FE	Y		Y		Y		Y	
Bidder FE	N		N		N		N	
Std-Dev. γ_{it}	0.35		0.33		0.36		0.33	
Std-Dev. η_{it}	0.09		0.09		0.11		0.09	
LLF/N	-2.407		-2.357		-2.761		-2.355	
Nb. Parameters	43		47		43		51	
Nb. Bids	2,737		2,737		682		2,737	
Nb. Invitations	10,210		10,210		2,638		10,210	
Nb. Firms	39		39		10		39	
Beliefs specification	Asymmetric		Asymmetric		Asymmetric		Asymmetric	

Table 5 presents the estimated parameters for the variable cost (top panel) and participation cost (bottom panel) functions. The first specification pools across bidder types, while the other three specifications account for differences in cost across bidders classified as competitive, collusive and naïve. In particular, specification (2) includes fixed-effects for bidder types, and specification (3) estimates the cost function only in the sample of competitive types. Accounting for bidder type fixed-effects leads to a meaningful improvement in fit, and we find that competitive types have lower marginal cost (≈ 300), and higher opportunity cost of participating. This is consistent with the interpretation that competitive bidders are more productive, and have are more likely to be constraint by their capacity. In specification (3), we also find that competitive types face a slightly higher disutility of distance. The results are inconsistent with respect to the effect of membership with the association on variable cost. This is because conditional on other characteristics, trucks affiliated with the association are more likely to be classified as "competitive". On the other hand, we find that trucks affiliated with the association are less likely to bid overall (across all specifications). We also find that single-truck bidder on average have higher cost, and are more likely to participate.

We use a mixture of normal with two mixtures in all specifications. The estimated standard deviation of γ ranges from 0.33 to 0.35 across specifications. The estimated distribution is a mixture of a normal distribution with a relatively small variance (85% weight on sd $= 0.28$) and a normal distribution with large variance (15% weight on sd $= 0.6$). This leads to a distribution of γ that exhibits more skewness than the normal distribution. The shape of the distribution is comparable across specifications.

The estimated distribution of the participation exhibits a relatively small variance. Recall that the standard-deviation of η_{it} corresponds to the inverse of the slope coefficient associated with the expected profit $\bar{\pi}_{it}$. We find that bidders are very sensitive to variation in the expected profits across auctions, leading to an estimate of σ_η that ranges from 0.09 to 0.11.

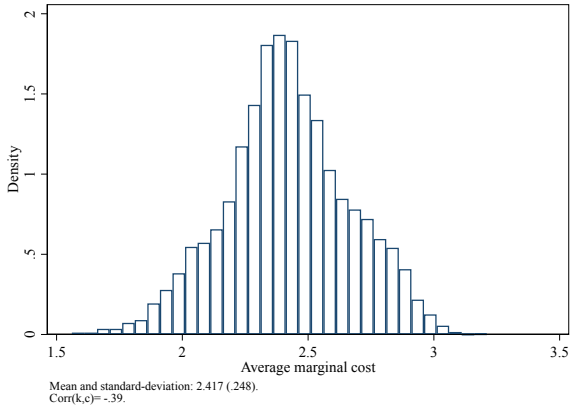
Finally, the last specification includes the number of invited bidders and the share of invited bidders belonging to the same garage as additional control variables. Under the null hypothesis of independent private value, the estimated cost distribution is independent of the number of bidders. The results are consistent with this hypothesis. The marginal effect of the number of invited bidders is close to zero both in the marginal cost and participation cost functions. We similarly include the fraction of invited bidders associated with bidder i's garage to test the hypothesis that firms submit independent bids. Once again, the point estimates are small and not statistically significant from zero. We therefore reject the hypothesis that bidders coordinate their bidding and participation decisions within garages.

Figure 9 presents the estimated distribution of average variable and fixed costs, as well as the distribution of expected profit for competitive and non-competitive types. The figures are generated using the estimates from Specification (1), which does not control directly for the type of bidders. In this specification, a bidder's competitive conduct is summarized its consideration-set $\mathcal{B}_i$, which in turn affects the expected profit of participating and the estimated marginal cost. The distributions are very similar when we in addition condition on the type of bidders directly in the cost function, as in Specifications 2 and 3.

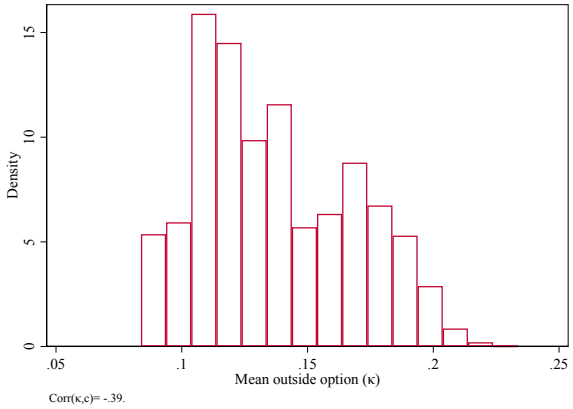
We estimate that the average marginal cost of serving clients is 24K, and is highly dispersed across bidders and clients. Most of that dispersion is due to observed differences distance between garages and clients, as well as the inclusion of a quarter-year fixed-effects which capture differences in cost

Figure 9: Distribution of expected costs and profits (Specification 1)

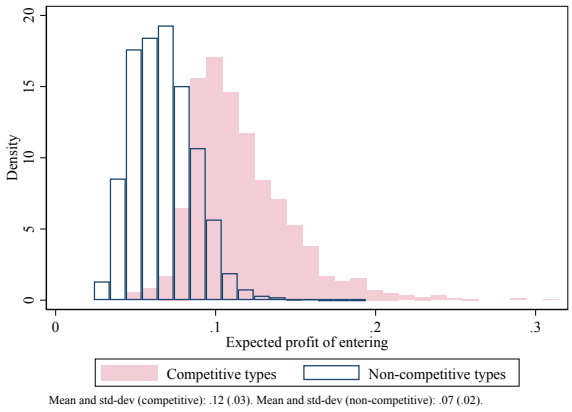
(a) Average marginal cost



(b) Average opportunity cost



(c) Expected profit of entering



across time. The average opportunity cost is similarly dispersed across firms. Recall that the mean of κ_{it} depends only on bidder characteristics and time, and reflects systematic differences in the value of working outside of the platform. The empirical distribution exhibits two modes due to the inclusion of indicator variables for single-truck bidders and for the largest garage (Foire). Importantly, for the counter-factual simulations, we estimate a negative correlation between the mean variable cost of providing the service and this outside option. In other words, firms with good outside options also tend to have lower marginal costs; for instance, because they are located on average further away from clients.

The distribution of expected profits in Figure 9c reflects a large amount of heterogeneity across bidders and clients. The main factors explaining this dispersion profitability are the proximity of the client to the bidder's garage, and the richness of the bidder's consideration set. Bidders with finer consideration-sets are more profitable because they are more likely to win and avoid ties, despite bidding lower amounts on average. We illustrate this point by plotting the estimated distribution of $\bar{\pi}_{it}$ for competitive and non-competitive types. Competitive bidders expect to earn 1,200 on average, compared to 700 for non-competitive types.

4 What is the effect of competition and platform design on prices and take-up?

4.1 A direct mechanism approach

This section maps the auctions modeled above into a mechanism design model, allowing us to abstract from the details of the marketplace to consider how the market fundamentals and the allocation rule determine the take-up of mechanized desludging on the platform. The model estimated above decomposes the firms' costs into "observable" and "unobservable" components, which allows us to study who are the most appropriate desludgers to invite to a given auction. We use a direct mechanism approach, rather than modeling the equilibrium outcome of a first-price sealed-bid auction, to simplify the computation of equilibrium outcomes under counter-factual invitation rules.

We consider an incentive-compatible direct mechanism in which a market designer chooses an invitation strategy auction t prior to observing the private-value components $(\gamma_{it}, \epsilon_{it})$ and consumer's willingness-to-pay (w_t). Similar to the current platform, we focus random invitation strategies. We analyze invitation lists that target invitations by proximity to the client and that vary the expected number of potential bidders. An invitation strategy for auction t, denoted by $\mathcal{A}_t$, is summarized by the number of invited bidders, and the probability of inviting each bidder $i = 1, \ldots, I$. We use A_t to denote the realized set of invited bidders in auction t, and $\tilde{A}_{it}$ the set of invited bidders who chose to participate conditional bidder i participating.

Conditional on $\mathcal{A}_t$, a direct mechanism for auction t is: (1) an entry strategy $a_i(\eta_{it}, \mathcal{A}_t)$ for each desludger $i \in A_t$ invited with positive probability, (2) the probability that a desludger i who enters is allocated the job, $P_{it}(\gamma_t, w_t, \tilde{A}_{it})$ given entry by the subset $\tilde{A}_{it} \subset A_t$, the household's realized value w_t, and the reported shocks of the desludgers γ_t, and (3) the payment that desludger i who entered receives, $T_{it}(\gamma_t, w_t, \tilde{A}_{it})$ given entry by the subset $\tilde{A}_{it} \subset \mathcal{A}_t$, the household's realized value w_t, and the reported shocks of the desludgers γ_{nt}. Thus, a direct mechanism is a tuple of functions

$$\mu_t = \left\{ a_i(\eta_{it}, \mathcal{A}_t), P_{it}(\gamma_{it}, \gamma_{-i,t}, w_t, \tilde{A}_{it}), T_{it}(\gamma_{it}, \gamma_{-it}, w_t, \tilde{A}_{it}) \right\}_{i \in A_t, \tilde{A}_{it} \subset A_t}. \tag{18}$$

The Revelation Principle guarantees that any auction or market has an equilibrium that is equivalent to the equilibrium of a direct mechanism in which the agents are induced to report their private information honestly. To replicate the timing of the real market, the platform suggests an entry strategy and bid function at the same time, so that the suggested bid function is independent of the realized entry of the other bidders. Thus, a direct mechanism is incentive compatible if each desludger i wishes to report his shock γ_{it} honestly during the bidding stage, and finds it in his best interests to enter in accordance with the suggestion $a_i(\eta_{it}, \mathcal{A}_t)$.

Consider the auction stage. In ignorance of the number of his competitors, an invited desludger i who entered then has the direct utility function

$$U_{it}(\hat{\gamma}_{it}, \gamma_{it}, \mathcal{A}_t) = \mathbb{E}_{\tilde{A}_{it}, \gamma_{-i,t}, w_t} \left[T_{it}(\hat{\gamma}_{it}, \gamma_{-i,t}, w_t, \tilde{A}_{it}) - P_n(\hat{\gamma}_{it}, \gamma_{-i,t}, w_t, \tilde{A}_{it})(\bar{c}_{it} + \gamma_{it}) \middle| \mathcal{A}_t \right]$$

and indirect utility function $U_{it}(\gamma_{it}, \mathcal{A}_t) = \max_{\hat{\gamma}} U_{it}(\hat{\gamma}_{it}, \gamma_{it})$. A direct mechanism is incentive compatible

if $U_{it}(\gamma_{it}, \gamma_{it}, \mathcal{A}_t) \geq U_{it}(\gamma'_{it}, \gamma_{it}, \mathcal{A}_t)$ for all γ_{it} and γ'_{it}, and individually rational if $U_{it}(\gamma_{it}, \mathcal{A}_t) \geq 0$. Note that the expectation operator integrates over the distribution of invited bidders A_t, and the distribution of participating $\tilde{A}_{it}$ bidders given A_t.

Theorem 1 (Myerson (1981)). *A mechanism μ_{it} is incentive compatible iff*

i. $\mathbb{E}_{\tilde{A}_{it}, \gamma_{-i,t}, w_t} \left[P_{it}(\hat{\gamma}_{it}, \gamma_{-i,t}, w_t, \tilde{A}_{it}) \middle| \mathcal{A}_t \right]$ *is decreasing in* $\hat{\gamma}_{it}$

ii.
$$\frac{\partial U_{it}(\gamma_{it}, \mathcal{A}_t)}{\partial \gamma_{it}} = -\mathbb{E}_{\tilde{A}_{it}, \gamma_{-i,t}, w_t} \left[P_{it}(\gamma_{it}, \gamma_{-i,t}, w_t, \tilde{A}_{it}) \middle| \mathcal{A}_t \right].$$

This implies that the expected payment to a desludger under any incentive compatible mechanism μ is

$$\mathbb{E}_{\tilde{\mathcal{A}}_t, \gamma_{nt}, w_t} \left[T_{it}(\gamma_{nt}, w_t, \tilde{A}_{it}) \middle| \mathcal{A}_t \right] = \mathbb{E}_{\tilde{A}_{it}, \gamma_{nt}, w_t} \left[P_{it}(\gamma_{nt}, w_t, \tilde{A}_{it}) \left\{ \bar{c}_{it} + \gamma_{it} + \frac{F(\gamma_{it})}{f(\gamma_{it})} \right\} \middle| \mathcal{A}_t \right].$$

The key piece of Theorem 1 is that it allows us to express the expected payment to the desludger as a function of the informationally adjusted cost of selecting desludger i to do the job: $\bar{c}_{it} + \psi_{it}(\gamma_{it})$. This decomposes the cost into three pieces. First, there is an observable part, $\bar{c}_{it}$, which function of distance and other covariates. The second part is the privately known shock, γ_{it}, which perturbs the desludger's true cost of the job. The third part is $F(\gamma_{it})/f(\gamma_{it})$, which is the desludger's informational rent: this represents the amount the platform must pay in excess of the true cost in order to induce entrants to reveal their private information honestly. The mean cost $\bar{c}_{it}$ plus the virtual cost $\psi(\gamma_{it})$ then represent the informationally adjusted cost to the platform, including incentive provision. Note that the information rent function is independent of i and t, because of our assumption that the distribution of the private-information shock is the same across bidders and clients, and that the auctioneer observed the characteristics of bidders and clients that determine $\bar{c}_{it}$.

We focus on mechanisms that allocate the job to the trucker with the lowest realized virtual cost. This leads to the following selection rule:

$$P_{it}(\gamma_{it}, \gamma_{i,-j}, w_i, \tilde{A}_{it}) = \begin{cases} 1 & \text{If } \bar{c}_{it} + \psi(\gamma_{it}) < \bar{c}_{it'} + \psi(\gamma_{it'}), \forall j' \in \tilde{A}_{it} \backslash j \text{ and } \bar{c}_{it} + \psi(\gamma_{it}) < w_i, \\ 0 & \text{Otherwise.} \end{cases} \tag{19}$$

This allows us the write the expected payment as function of the auction winning probability and the demand function from consumers:

$$\begin{aligned}
&E_{\tilde{A},\gamma,w} \left[P_{it}\left(\gamma_{it}, \gamma_{i,-j}, w_t, \tilde{A}\right) \times \left[\bar{c}_{it} + \psi(\gamma_{it})\right] \middle| \mathcal{A}_t \right] \\
&= E_{\tilde{A}} \left[\int \underbrace{D_t(\bar{c}_{it} + \psi(\gamma_{it}))}_{\Pr(w_t > \bar{c}_{it} + \psi(\gamma_{it}))} \prod_{k \in \tilde{A}_{-i}} \underbrace{\left[1 - F\left(\psi^{-1}(\bar{c}_{it} + \psi(\gamma_{it}) - \bar{c}_{kt})\right)\right]}_{\Pr(\bar{c}_{it} + \psi(\gamma_{it}) < \bar{c}_{ik} + \psi(\gamma_{ik}))} \times \left[\bar{c}_{it} + \psi(\gamma_{it})\right] f(\gamma_{it}) d\gamma_{it} \middle| \mathcal{A}_t \right] \\
&= E_{\tilde{A}_{-i}} \left[\int G_{it}(\gamma_{it}, \tilde{A}_{-i}) \times \left(\bar{c}_{it} + \psi(\gamma_{it})\right) f(\gamma_{it}) d\gamma_{it} \right]
\end{aligned}$$

where $\tilde{A}_{-i}$ denotes a random participant list excluding bidder i, and $\psi^{-1}(\cdot)$ is an inverse-mapping $b_{it} =$

$\psi(\gamma_i) \to \gamma_{it} = \psi^{-1}(b_{it})$. As before, the function $G_{it}(\gamma_{it}, \tilde{A}_t)$ denotes the transaction probability of firm i in auction t, conditional on (random) participation list $\tilde{A}_t$, and idiosyncratic cost realization γ_{it}.

Theorem 1 also allows us to compute the desludgers' expected profits of entering the auction for any invitation list strategy:

$$\begin{aligned}
\bar{\pi}_{it}(\mathcal{A}_t) &= E_{\tilde{A},\gamma_{it},\gamma_{-i,t},w_t}\left[\int_{\gamma_{it}}^{\infty} P_{it}(z,\gamma_{-i,t},w_t,\tilde{A})dz\middle|\mathcal{A}_t\right] \\
&= E_{\tilde{A}}\left[\int\left\{\int_{\gamma_{it}}^{\infty} D_t(\bar{c}_{it}+\psi(z))\prod_{k\in\tilde{A}\setminus j}\Pr(\psi(\gamma_{kt})>(\bar{c}_{it}-\bar{c}_{kt})+\psi(z))dz\right\}f(\gamma_{it})d\gamma_{it}\middle|\mathcal{A}_t\right] \\
&= E_{\tilde{A}_{-i}}\left[\int\int_{\gamma_{it}}^{\infty} G_{it}(z,\tilde{A}_{-i})dz f(\gamma_{it})d\gamma_{it}\middle|\mathcal{A}_t\right] = E_{\tilde{A}_{-i}}\left[\Pi_{it}\left(\tilde{A}_{-i}\right)|\mathcal{A}_t\right]
\end{aligned}$$

Let $\rho_{it}(\mathcal{A}_t)$ denotes the entry probability of bidder i in auction t. The equilibrium beliefs if bidder i regarding the distribution of $\tilde{A}$ is given the product of N_t binomial probabilities given:

$$h(\tilde{A}_{-i}|A_t,\rho_t) = \begin{cases} \prod_{k\in\tilde{A}_{-i,t}}\rho_{kt}(\mathcal{A}_t)^{\tilde{a}_{kt}}\left(1-\rho_{kt}(\mathcal{A}_t)\right)^{1-\tilde{a}_{kt}} & \text{If } \tilde{A}_{-i,t}\subset A_{-i,t}, \\ 0 & \text{Otherwise.} \end{cases} \tag{20}$$

where $\tilde{a}_{kt}$ is element k of $\tilde{A}_{it}$ (i.e. indicator variable equal to one if k participates).

The expected profit given belief ρ can therefore be re-written as:

$$\bar{\pi}^{\rho}_{it}(\mathcal{A}_t) = \sum_{A_{-i,t}}\mathcal{P}\left(A_{-i,t}|\mathcal{A}_t\right)\left[\sum_{\tilde{A}_{-i}}h(\tilde{A}_{-i}|A_{-i,t},\rho_t)\Pi_{it}\left(\tilde{A}_{-i}\right)\right] \tag{21}$$

where $\mathcal{P}\left(A_{-i,t}|\mathcal{A}_t\right)$ is the probability distribution of invitation lists given invitation rule $\mathcal{A}_t$.

A Perfect-Bayesian Nash equilibrium is defined as a threshold entry rule $\bar{\eta}^{\rho}_{it}(\mathcal{A}_t)$ satisfying the following fixed-point solution:

$$\begin{aligned}
\bar{\eta}^{\rho}_{it}(\mathcal{A}_t) &= \bar{\pi}^{\rho}_{it}(\mathcal{A}_t) - \kappa_{it} \qquad (22)\\
\text{Such that,} \qquad \rho_{it}(\mathcal{A}_t) &= \Phi\left(\frac{\bar{\eta}^{\rho}_{it}(\mathcal{A}_t)}{\sigma_\eta}\right) \quad \forall i\in\mathcal{A}_t \text{ and } t=1,\ldots,T.
\end{aligned}$$

With these results, the welfare delivered by *any* incentive compatible rule μ must equal

$$\mathbb{E}_{\tilde{A}_t,w_t,\gamma_t}\left[\sum_{i\in\tilde{A}_t}P_{it}(\gamma_{nt},w_t,\tilde{A}_t)\left(w_t-\bar{c}_{it}-\psi(\gamma_{it})\right)\middle|\mathcal{A}_t\right], \tag{23}$$

given that $\mathbb{E}_{\tilde{A}_t,\gamma_{-i,t},w_t}[P_{int}(\hat{\gamma}_{it},\gamma_{-i,t},w_t,\tilde{A}_t)|\mathcal{A}_t]$ non-increasing in $\hat{\gamma}_{it}$.

Note that this expression of welfare applies, basically, for any invitation scheme and any allocation rule in which reporting a lower cost perturbation γ_{it} results in a higher probability of winning. This allows us to determine welfare under alternative invitation rules, making it feasible to evaluate the social benefits of different platforms.

This formulation captures the fundamental trade-offs of the environment. By selecting a more com-

petitive A_t, the expected profits of the invitees fall, and they are less likely to enter. The welfare of the household and efficiency of the market must be traded off with the rents left on the table for the desludgers. Without balancing the motives of both sides of the market, neither side will find the result attractive, and the market can unravel.

Solution algorithm

To solve the model, we must find the equilibrium entry probability vector associated with each counterfactual stochastic invitation rule $\mathcal{A}_t$. There are two main numerical difficulties associated with this problem. First, the fact that firms are heterogenous both in terms of their outside option and average variable cost (i.e. $\bar{c}_{it}$ and κ_{it}) implies that the expected profit conditional of a given participation list $\tilde{A}$ requires to compute the product over rival bidders' CDFs, and invert their virtual cost functions. Second, the number of potential lists $\tilde{A}$ increases exponentially with the number of firms invited to bid (curse of dimensionality).

To alleviate these two problems, first note that the expected profit conditional on the realization of participation list $\tilde{A}_{-i}$ is independent of ρ:

$$\begin{aligned}
\Pi_{it}\left(\tilde{A}_{-i}\right) &= \int\int_{\gamma_{it}}^{\infty} G_{it}(z, \tilde{A}_{-i}) dz f(\gamma_{it}) d\gamma_{it} \\
&= \int\int_{\gamma_{it}}^{\infty} D_t(\bar{c}_{it} + \psi(z)) \prod_{k \in \tilde{A}_{-i}} \left[1 - F\left(\psi^{-1}(\bar{c}_{it} + \psi(z) - \bar{c}_{kt})\right)\right] dz f(\gamma_{it}) d\gamma_{it} \\
&\approx \sum_{k=1}^{K} D_t(\bar{c}_{it} + \psi(x_{k1})) \prod_{k \in \tilde{A}_{-i}} \left[1 - F\left(\psi^{-1}(\bar{c}_{it} + \psi(x_{k1}) - \bar{c}_{kt})\right)\right] f(x_{k2}) \omega_k \qquad (24)
\end{aligned}$$

where $\{x_{k1}, x_{k2}, \omega_k\}_{k=1,\ldots,n}$ is a matrix of gaussian quadrature integration grids and weights, such that $x_{k1} \geq x_{k2}$ for all k. Equation (24) can be "pre-calculated" for a large number of lists $\tilde{A}$, and store in memory. Note that the inverse virtual cost function can be approximated using a smooth monotonic function. At the estimated parameters, we find that the distribution of idiosyncratic cost implies that the inverse-mills ratio is a monotonically increasing of bidders' types.[12]

To address the curse of dimensionality problem in the number of potential participation lists, we use Monte-Carlo simulation methods. Importance sampling is particularly well suited here, because it can take advantage of the fact that the profit function is dependent of the participation probability (e.g. **? ?**). For each auction t, we simulate S invitation lists A_s that are consistent with $\tilde{A}_t$, and S participation lists $\tilde{A}_s$ that are consistent with an estimate of the entry probability of players. We use the observed entry probability of each player to construct this set. We then solve the game by iterating over the entry probability assuming that players form expectations about $\tilde{A}_t$ by considering only the (fixed) simulated set of lists $\{\tilde{A}_t\}_{s=1,\ldots,S}$. Here is a sketch of the simulation algorithm:

1. Generate $S \to \infty$ random invitation list using the number of invited bidders and sampling probability given by $\mathcal{A}_t$.

[12] We approximation of the inverse function using a piece-wise local polynomial regression of degree 3 with 5 discrete segments.

2. Generate $S \to \infty$ participation lists $\tilde{A}_{-i,t}$ by sampling participation decisions from an initial estimate of the participation probability ρ_{it}^0 conditional on A_t.

3. For each simulated participation lists, calculate the expected profit of participating: $\Pi_{it}(\tilde{A}_{-i})$.

4. Iteration $l \geq 1$

   For $i = 1, \ldots, N$:

   2.i Compute belief of bidder i:

$$h(\tilde{A}_{-i,s}|\mathcal{A}_t, \rho_t^l) = \prod_{k \in A_{-i,t}} \tilde{\rho}_{kt}^l(\mathcal{A}_t)^{\tilde{a}_{kt}} \left(1 - \tilde{\rho}_{kt}^l(\mathcal{A}_t)\right)^{1-\tilde{a}_{kt}}, \quad \forall s = 1, \ldots, S$$

   where $\tilde{\rho}_{kt}^l(\mathcal{A}_t) = \rho_{kt}^{l-1}(\mathcal{A}_t)$ if $k > i$, and $\tilde{\rho}_{kt}^l(\mathcal{A}_t) = \rho_{kt}^l(\mathcal{A}_t)$ if $k < i$ (i.e. Gauss-Siedel updating). Use the belief probability distribution to compute the expected profit of participating:

$$\hat{\pi}_{it}^{\rho^l}(\mathcal{A}_t) = \frac{h(\tilde{A}_{-i,s}|\mathcal{A}_t, \rho_t^l)}{\sum_{s'=1}^{S} h(\tilde{A}_{-i,s'}|\mathcal{A}_t, \rho_t^l)} \Pi_{it}(\tilde{A}_{-i})$$

   Update the participation probability of player i:

$$\rho_{it}^l(\mathcal{A}_t) = \Phi\left(\frac{\hat{\pi}_{it}^{\rho^l}(\mathcal{A}_t) - \kappa_{it}}{\sigma_\epsilon}\right)$$

3. Repeat steps 4.1-4.N until convergence: $||\rho_{it}^l(\mathcal{A}_t) - \rho_{it}^{l-1}(\mathcal{A}_t)|| < \eta$.

To solve the model we set the number of simulated lists $S = 5{,}000$, and sample initial participations using the average participation frequency of each bidder (i.e. $\rho_{it}^0 = \hat{\rho}_i$ for all t). We consider five alternative invitation rules:

- **Uniform:** Uniform invitation probabilities with an expected number of bidders equal to 10
- **Distance weight:** Inverse distance weight invitation probabilities with fixed number of invited bidders equal to 10, 15 and 20[13]
- **Full:** All 39 active bidders invited to bid on every auctions

We then simulate auction outcomes using the equilibrium entry probability of each invited bidder. We do so by simulating 5,000 entry decisions and private-value costs for each simulated participation/invitation lists.

4.2 Counter-factual results

We use the abovementioned procedure to compute counterfactual equilibrium outcomes under alternative platform designs. Our objective is to quantify the main market frictions determining the platform's

[13]We define invitation probabilities that are proportional to the square inverse distance between a bidder garage and the client.

performance: imperfect competition and asymmetric information. First, we simulate offers and take-up rates using an invitation rule that mimics the one used in the observed platform. We use this experiment to quantify the effect of eliminating non-competitive bidding strategies on the platform. Second, we compare the same outcomes in counter-factual environments with increasing levels of competition to discuss the effect of improving the design of the auction platform on prices and take-up. Finally, we quantify the welfare cost of asymmetric information and strategic entry by comparing the expected consumer surplus under second-best and first-best allocations and the second-best allocation without participation costs.

Eliminating non-competitive conduct

Table 6: Observed vs counter-factual auction outcomes with uniform invitation probabilities lists

(a) Specification 1: All bidders

	Dist. (km)	1(Bid)		Lowest bid		Price		1(Trans.)	
	(1)	(2) Obs.	(3) CF	(4) Obs.	(5) CF	(6) Obs.	(7) CF	(8) Obs.	(9) CF
Downtown	16.96	0.22	0.42	2.32	2.15	2.22	1.95	0.39	0.49
Guediawaye	14.96	0.19	0.48	2.43	2.23	2.36	2.03	0.22	0.37
Niayes	21.91	0.16	0.40	2.78	2.63	2.69	2.33	0.15	0.26
Pikine	11.03	0.21	0.45	2.24	2.11	2.00	1.88	0.24	0.34
Thiaroye	18.37	0.20	0.50	2.46	2.20	2.39	2.02	0.32	0.41
Total	18.02	0.19	0.45	2.51	2.32	2.37	2.10	0.26	0.36

Number of simulations: 5,000 auctions and 5,000 invitation lists.
Transaction prices are expressed in 10,000 CFA.
Invitation list: Uniform probability weight (10 firms).

(b) Specification 3: Competitive types

	Dist. (km)	1(Bid)		Lowest bid		Price		1(Trans.)	
	(1)	(2) Obs.	(3) CF	(4) Obs.	(5) CF	(6) Obs.	(7) CF	(8) Obs.	(9) CF
Downtown	16.96	0.22	0.28	2.32	1.76	2.22	1.51	0.39	0.62
Guediawaye	14.96	0.19	0.26	2.43	2.24	2.36	1.90	0.22	0.38
Niayes	21.91	0.16	0.19	2.78	2.79	2.69	2.28	0.15	0.23
Pikine	11.03	0.21	0.23	2.24	2.23	2.00	1.82	0.24	0.32
Thiaroye	18.37	0.20	0.27	2.46	2.30	2.39	1.97	0.32	0.39
Total	18.02	0.19	0.24	2.51	2.39	2.37	2.00	0.26	0.35

Number of simulations: 5,000 auctions and 5,000 invitation lists.
Transaction prices are expressed in 10,000 CFA.
Invitation list: Uniform probability weight (10 firms).

Tables 6a and 6b presents the results of the first exercise. The first panel uses the cost estimates from Specification (1) which uses bidding strategy from all bidders, and assume that firms have the same cost

function across all conduct types. The second panel uses the cost estimates from Specification (3) which estimates the cost function only using the sample of competitive types. The simulation results using estimates from Specification (2) are omitted, and are very similar to the ones obtained from Specification (1). The main benefit of using Specification (3) to estimate marginal cost and participation cost is that the profit maximization assumption is more likely to hold in the sample of competitive bidders. The downside however, is that if competitive types have different cost functions, the parameter estimates might not be representative of the population of truckers.

Focussing on Table 6a, we see that eliminating non-competitive conduct on the platform can have meaningful impact on the distribution of offers and demand. Recall that the invitation list are constructed to mimic the invitation list used on the platform. The number of invited bidders is random, and on average the simulation invites 10 bidders per auction. This number is smaller than the actual number of invited bidders on the platform, but matches the average number of bidders invited among those who submitted at least one bid. One of the largest change is the increase in the participation rate from 19% to 45%. This mostly due to the fact that most bidders use very coarse bidding strategy, which leads to low expected profit. The direct mechanism uses continuous bids and allocate the contract to the lowest cost supplier, which increases significantly the number of bids received. Eliminating focal prices also contributes to lowering the average lowest bid received to consumers from 25K to 23.2K. Since consumers are very price elastic, this leads to a 13% decrease in price for consumers, and a 35% increase in demand (from 27% to 36% take-up). Importantly, we predict that eliminating non-competitive conduct improves take-up across all neighborhoods. This is particularly true in neighborhoods that are located close to truckers (e.g. Pikine).

Table 6b shows that most of these predictions are robust to our estimates of cost. Using the cost estimates from Specification (3), we find that the take-up rate would increase from 27% to 35% from eliminating non-competitive conduct, and transaction prices would decrease by a slightly larger margin (from 24K to 20K). The mechanism for this increase in competition is slightly different across these two specifications. The cost estimates from the sample of competitive bidders is associated with lower marginal cost, but higher opportunity cost of bidding. As a result, participation on the platform is predicted to increase only by 5 percentage points, which leads to fewer bids and a higher probability of receiving zero bids. The failure rate is nearly zero in the top panel, but about 9% in the bottom panel. In contrast, 17% of auctions in the observed platform resulted in zero bids. These differences explain why the average accepted offer is nearly 5% lower in Table 6b compared to Table 6a, and yet the transaction probability is slightly lower.

In Houde et al. (2024) we constructed a similar counter-factual exercise by simulating the outcome of auctions in which all active bidders used the competitive-type policy function. This exercise ignores the equilibrium response of bidders from a market-wide change in conduct. In contrast, in our current simulation results, the entry probability of invited bidders is allowed to respond endogenously to the change in conduct. Although the simulation exercise in that paper was conducted on a slightly different sample, the results are broadly consistent. The partial equilibrium simulation predicted a 23% increase in take-up, compared to 35% for the equilibrium model. Similarly, accounting for equilibrium response leads to more pronounced decrease in prices for consumers: -9.3% vs 13% or 15% here.

Competitive invitation lists

Table 7: Auction outcomes across different invitation lists

(a) Specification 1: All bidders

	Uniform	Inverse distance weights			Full list
	(1)	(2)	(3)	(4)	(5)
	$\bar{N} = 10$	$N = 10$	$N = 15$	$N = 20$	$N = 39$
	Mean	Mean	Mean	Mean	Mean
Total distance (km)	17.121	15.349	15.516	15.676	17.120
1(Transaction)	0.359	0.371	0.401	0.420	0.457
Transaction price	2.096	2.073	2.014	1.975	1.896
Winning bid	2.322	2.271	2.165	2.107	2.005
1(Zero bids)	0.007	0.004	0.001	0.000	0.000
1(Participation)	0.454	0.466	0.405	0.364	0.275
Observations	2636	2636	2636	2636	2636

Simulation: 5,000 auctions and 5,000 invitation lists.

Prices are expressed in 10,000 CFA.

(b) Specification 3: Competitive types

	Uniform	Inverse distance weights			Full list
	(1)	(2)	(3)	(4)	(5)
	$\bar{N} = 10$	$N = 10$	$N = 15$	$N = 20$	$N = 39$
	Mean	Mean	Mean	Mean	Mean
Total distance (km)	17.121	15.349	15.516	15.676	17.120
1(Transaction)	0.349	0.359	0.400	0.425	0.467
Transaction price	2.005	1.986	1.939	1.906	1.841
Winning bid	2.395	2.364	2.229	2.144	2.010
1(Zero bids)	0.093	0.082	0.043	0.026	0.012
1(Participation)	0.240	0.244	0.211	0.189	0.137
Observations	2636	2636	2636	2636	2636

Simulation: 5,000 auctions and 5,000 invitation lists.

Prices are expressed in 10,000 CFA.

Tables 7a and 7b present the results of the second exercise in which we vary competition on the platform. The first column corresponds to the uniform invitation list discussed above. Columns (2-4) use an inverse-distance invitation probability weight, and column (5) invites all 39 active bidders. Each entry corresponds to the average auction outcomes across all neighborhoods, and the top row reports the average distance between clients and invited bidders.

The comparison between columns (1) and (2) measures effect of targeting invitations by distance (holding fixed the number of invited bidders), which reduces the average distance on the platform from 17.121 KM to 15.349 KM. Both specifications lead to similar conclusions. Reducing the average distance leads to more competitive offers, but the effects are relatively small in magnitude. In the top panel, we find that the average winning bid decreases by 2.2%, and the average accepted offer decreases by 1%. This leads to 1 percentage point increase in the take-up rate. The main reason for this is the relatively

small improvement in the proximity to clients. Targeting invitations leads to a roughly 2 KM decrease in distance, which corresponds to a 380 CFA decrease in variable cost (specification 1). The cost decrease is large under specification (3), but the fact that firms are less likely to participate offset this improvement in matching. The limited effect invitation targeting highlights the fact that many clients are located relatively far from the main garages, which reduces the potential benefit of improving the matching between clients and truckers.

The rest of the table illustrates that increasing the number of invited bidders on the platform leads to much large price reductions. Increase the number of potential bidders from 10 to 15 increase the take-up rate by 8% in Specification 1 and 11.4% in Specification 4. According to the estimates from Specification 1, this increase in take-up is mostly due to a price reduction, since the number of failed auctions is nearly zero across all specifications. In the bottom panel, the effect is more nuanced, since inviting more bidders leads to a large reduction in the number of auctions with zero bids, as well as a reduction in the cost of participating bidders.

Both specifications suggest that it is optimal for the platform to invite more bidders in each auction. In particular we find that despite the fact that the participation rate falls with the number of invited bidders (due to a reduction in expected profit), the average lowest bid is strictly decreasing in the number of invited bidders. We find that using a full list of bidders would lead to a take-up rate of 46%, compared to 27% in the observed platform. Transaction prices obtained from the platform would also fall to 18.9K and 18.4K, in Specification 1 and 3 respectively. Although those prices are higher than those observed in the Rufisque Arrondissement (15K), the cost of serving clients on the platform is likely much higher due to the fact that clients are located further away from the main garages.

Increasing competition on the platform is beneficial because bidders differ in the cost of providing the service and in their opportunity cost of submitting a bid. Because a large portion of these cost differences are privately known to bidders, it is beneficial for the platform to invite a large number of bidders. This conclusion could be different if the importance of private information on the platform was reduced. For instance, if firms had a small variance of γ_{it} the decrease in participation could outweigh the effect of competition on offers. In this case the platform could lead to higher prices when the platform invites all bidders, and it would be beneficial to invite an intermediate number of bidders.

Consumer welfare decomposition

Finally, Tables 8a and 8b decompose the losses in consumer-surplus from private information and strategic entry compared to the hypothetical first-best outcome that corresponds to perfect competition, where all desludgers bid their true costs. As before, we compare the model predictions obtained using the cost estimates from Specification 1 (all bidders cost) and Specification 3 (competitive types cost). The first-best allocation assumes that all 39 invited bidders participate, and the price column measures the average accepted competitive offer. This leads to a large price decrease relative to the second-best outcome across all neighborhoods, especially in Arrondissements located close to garages and treatment centers (e.g., Pikine). Eliminating both frictions increases consumer surplus by 88%, and increases the take-up rate by nearly 50%. In the first-best, roughly 70% of clients accept the platform's offer.

The middle panel isolates the welfare cost due to participation. Relative to the second-best allocation, the model predicts that eliminating the participation costs would decrease average prices from

Table 8: Welfare cost of private information and strategic entry in the platform

(a) Specification 1: All bidders

	2nd Best			2nd Best w/o Entry			1st Best		
	(1) Price	(2) Trans.	(3) Surplus	(4) Price	(5) Trans.	(6) Surplus	(7) Price	(8) Trans.	(9) Surplus
Downtown	1.714	0.614	0.586	1.487	0.704	0.744	1.251	0.786	0.932
Guediawaye	1.837	0.466	0.365	1.543	0.595	0.534	1.314	0.693	0.695
Niayes	2.132	0.356	0.249	1.810	0.497	0.400	1.587	0.599	0.536
Pikine	1.689	0.443	0.337	1.390	0.575	0.503	1.162	0.674	0.659
Thiaroye	1.820	0.515	0.426	1.544	0.632	0.597	1.313	0.726	0.767
Total	1.896	0.457	0.363	1.603	0.583	0.526	1.375	0.680	0.684

Number of simulations: 5,000 auctions and 5,000 invitation lists.

Transaction prices are expressed in 10,000 CFA.

Invitation list: Full (39).

(b) Specification 3: Competitive bidders

	2nd Best			2nd Best w/o Entry			1st Best		
	(1) Price	(2) Trans.	(3) Surplus	(4) Price	(5) Trans.	(6) Surplus	(7) Price	(8) Trans.	(9) Surplus
Downtown	1.280	0.766	0.899	1.019	0.854	1.130	0.846	0.895	1.287
Guediawaye	1.730	0.504	0.416	1.371	0.682	0.663	1.234	0.738	0.766
Niayes	2.132	0.339	0.240	1.725	0.554	0.463	1.590	0.616	0.548
Pikine	1.661	0.441	0.341	1.299	0.632	0.573	1.159	0.692	0.673
Thiaroye	1.804	0.513	0.431	1.450	0.686	0.675	1.312	0.741	0.780
Total	1.841	0.467	0.391	1.475	0.650	0.627	1.336	0.707	0.728

Number of simulations: 5,000 auctions and 5,000 invitation lists.

Transaction prices are expressed in 10,000 CFA.

Invitation list: Full (39).

Table 9: Consumer welfare and take-up rate for counter-factual platform designs

	Spec. 1: All bidders		Spec. 3: Comp. type	
	(1)	(2)	(3)	(4)
	Consumer surplus	1(Transaction)	Consumer surplus	1(Transaction)
Observed platform	0.150	0.246	0.150	0.246
Second-best: $\bar{N} = 10$	0.254	0.351	0.267	0.343
Second-best: $N = 39$	0.355	0.449	0.384	0.460
Second-best (w/o entry)	0.517	0.575	0.618	0.644
First-best	0.672	0.672	0.718	0.701

Number of simulations: 5,000 auctions and 5,000 invitation lists.

Average consumer surplus are expressed in 10,000 CFA.

CFA 18,960 to 16,030 (15.5%) and increase take-up from 45.7% to 58.3%. This would increase consumer-surplus by 45% (or CFA 1,630). This corresponds to the welfare cost of strategic entry in the platform. Recall that inviting all 39 bidders decreases participation probability from 45.5% to 27.5% in Specification 1. This reduction in competition on the platform explains slightly more than half of the consumer surplus difference between the second and first best allocation. Therefore, we estimate that private information accounts for 49% of the welfare loss relative to first-best, while strategic entry and private information together reduce welfare by 88%, from 6,840 to 3,630.

The results Specification 3, summarized in Table 8b, paint a similar picture. Both estimated cost functions lead to a similar welfare cost of private information and strategic entry. The first-best allocation is associated with a 86% increase in consumer surplus. Roughly 70% of this increase is due to the elimination of participation cost, and 30% due to private information about γ. This is because the model estimated using only competitive types is associated with lower participation, and lower variable cost.

Not surprisingly, the elimination of market frictions leads to even larger welfare gains when compared with the current platform. Table 9 compares the average consumer surplus and take-up rate in the observed platform with the counter-factual platform designs under increasing levels of competition. The current platform generates an average consumer surplus of 1,500 CFA and a transaction probability of 25%. The first-best allocation would increase consumer welfare by 350% in Specification 1 (and 375% in Specification 3). In both specifications, eliminating non-competitive conduct under the current uniform invitation probability rule increases consumer surplus by 1,000 CFA, and the transaction probability by ten percentage points. Improving the design of the platform by inviting all bidders more than doubles the consumer surplus generated by the platform and increases the take-up rate to 45%. Therefore, we estimate that roughly half of the improvement in consumer welfare obtained from optimally designing the platform is obtained by increasing the number of invited bidders, while the other half is obtained by eliminating non-competitive conduct associated with coarse bidding.

5 Conclusion

To be completed...

References

Asker, John. 2010. "A Study of the Internal Organization of a Bidding Cartel." *American Economic Review* 100:724–762.

Beaman, Lori, Jeremy Magruder, and Jonathan Robinson. 2014. "Minding small change among small firms in Kenya." *Journal of Development Economics* 108:69–86.

Coppejans, Mark. 2001. "Estimation of the binary response model using a mixture of distributions estimator (MOD)." *Journal of Econometrics* 102:231–269.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica* 68 (3):525–574.

Hortaçsu, Ali and Jakub Kastl. 2012. "Valuing dealers' informational advantage: A study of Canadian treasury auctions." *Econometrica* 80 (6):2511–2542.

Hortaçsu, Ali and David McAdams. 2010. "Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market." *Journal of Political Economy* 118 (5):833–865.

Houde, Jean-François, Terrence Johnson, Molly Limscomb, and Laura Schechter. 2024. "Imperfect Competition and Sanitation: Evidence from randomized auctions in Senegal." NBER working paper no. 30514.

Klein, ad Spady Richard, Roger. 1993. "An Efficient Semiparametric Estimator for Binary Response Models." *Econometrica* 61 (2).

Krasnokutskaya, Elena and Katja Seim. 2011. "Bid Preference Programs and Participation in Highway Procurement Auctions." *American Economic Review* 101 (6):2653–86.

Myerson, R. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1):58–73.

Rivers, D. and Q.H Vuong. 1988. "Limited Information Estimators and Exogeneity Tests for Simultaneous Probit Models." *Journal of Econometrics* 39:347–366.